Supporting Information

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

by Rick Quax, Drona Kandhai, and Peter M.A. Sloot

Table of contents

S1 Background of information dissipation

S1.1 Information theory

The amount of information that is stored in a variable is the minimum number of yes/no questions that is needed to determine a value for the variable[1]. The value of a variable that encodes the result of a fair 50%-50% coin toss can be uniquely identified by at least one yes/no question, namely 'did the coin toss result in heads?' We say that it stores 1 *bit* of information. More generally, a variable with N equally probable values stores $\log_2 N$ bits. This is the maximum amount of information that a variable with N possible values can store.

The information stored in a variable can be less than $\log_2 N$ if its values have different probabilities. Suppose, for instance, that we toss a coin of which we know it results in heads 90% of the time. Intuitively, the outcome of a toss is less informative because we already anticipate it in part. In the extreme case of a coin with two identical sides the outcome of a toss provides zero information because the question of its outcome is already completely answered beforehand.

In general, the number of bits that is stored in a variable s with possible outcomes $\{v_1, v_2, \ldots, v_n\}$ is the Shannon entropy

$$H(s) = -\sum_i p_i \log p_i ,$$

where p_i is the chance that the value of s is v_i. This quantity has two meanings. Firstly it is the number of bits that one must obtain, through measurement or inference, in order to identify the outcome of s, historically interpreted as the 'missing information'. The second meaning is that s is capable of storing $H(s)$ bits of information about other variables, which we explain next.

Let us interpret variable s_1 to encode the state of one dynamical unit and variable s_2 to encode the state of another dynamical unit. We refer to s_1 and s_2 as states and their values as instances. According to the second meaning of Shannon entropy, learning the instance of s_1 can provide between zero and $H(s_1)$ bits of information about the instance of s_2. It is non-zero in case the states s_1 and s_2 are correlated or cause-and-consequence, such that the fact that s_1 is in a particular instance tells us something about the instance of s_2. If s_1 and s_2 are independent processes then this *mutual information* is zero. If, on the other hand, s_2 encodes for instance the state of a tossed coin that tends to be equal to the state of another tossed coin s_1, then the more s_2 depends on s_1 the more information about s_1 is stored in s_2 (and vice versa).

The amount of this mutual information is

$$I(s_1 \mid s_2) = H(s_1) - H(s_1 \mid s_2), \quad (1)$$

where $H(s_1 \mid s_2)$ is the conditional variant of $H(s)$. In words, knowing s_2 reduces the number of unknown bits about the outcome of s_1 from $H(s)$ to $H(s_1 \mid s_2)$. We can interpret $I(s_1 \mid s_2)$ as the amount of information that the variable s_2 stores about the variable s_1.

Information can be transferred between interacting states. Here, an interaction between two states means that one state (partly) depends on the other state, vice versa, or both. Suppose that two interacting dynamical units s_1 and s_2 form a system, and suppose that both units are equally influenced by other factors outside this system. The information that is stored by the state s_1 about the other state s_2 now consists of two parts: an amount I_{corr} which both states have in common because they are subject to the same external influence (creating a correlation), and an additional I_{int} which is due to the interaction between the states (creating additional correlation)[2]. The information I_{int} is present in the variable s_1 at first and not in s_2 in case there is no interaction, so we can say that information is transferred, by means of interactions, from one state to another. Although I_{corr} did not transfer directly through the interaction between s_1 and s_2 it did transfer through the external interactions in a similar way to both s_1 and s_2. The reason that it creates mutual information between s_1 and s_2 is because both states store the same I_{corr} bits of information, which makes the stored information in both states overlap with each other.

S1.2 Information dissipation

In a system of interacting units we can let s_i^t correspond to the state of unit i at time t. The more random the interaction among the connected units, the less information is transmitted between the states and therefore the lower the correlation between the states of the units. These connections form a network in which information about a unit s is transmitted to its neighbors, which store it in their state. Each neighbor subsequently transmits information about its own state, which partly consists of the information from unit s, and so on, inducing the percolation and mixing of many different pieces of information through the system.

The distance that information travels is a direct measure of how far the dynamics of a unit percolates through the system. If there are no interactions then all information is stationary, and the units are incapable of system-wide state transitions. In the other extreme, if each unit would copy the state of another unit, then information can travel infinitely far and even a single unit can cause the entire system to change its state. For real systems, this *information dissipation length* (IDL) lies somewhere in between these two extremes. Since critical phenomena require a synchronous change of state of a majority of units in the system, our hypothesis is that we can quantify the plausibility of self-organized critical phenomena by a significant increase of IDL.

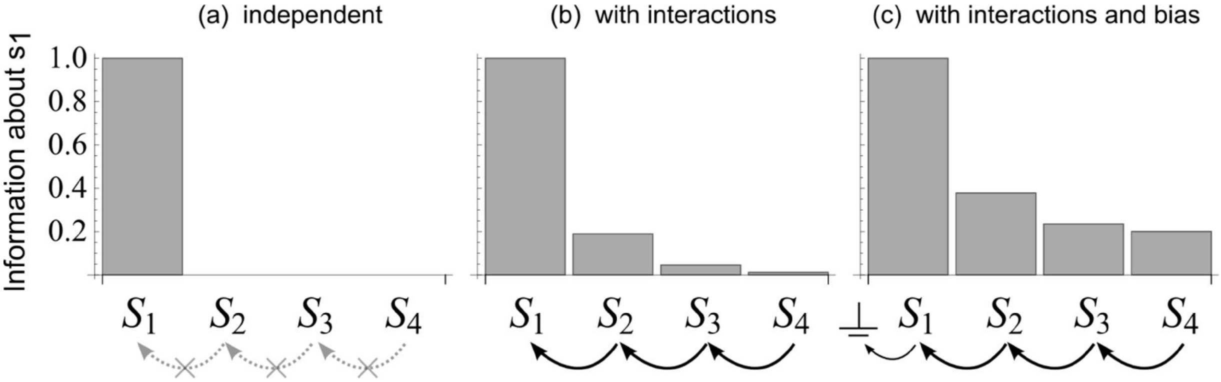

Figure S1: Three different modes of information transfer in a one-dimensional sequence of coin tosses. (a) If a subsequent coin toss does not depend on the previous outcome, then information remains local and is not transferred. (b) If each coin toss tends to be equal to the previous outcome, then information about the first outcome transfers and dissipates to zero. (c) If each coin toss is biased then each outcome provides the transferred information in addition to the constant prior information due to the bias. The dissipation time of information is not affected by the presence of prior information.

S1.3 Calculating the information dissipation length in IRS prices

We first illustrate how information dissipates through a one-dimensional system by the example of coin flips. Then we present an algorithm that can be used to measure the IDL in the prices of IRS of different maturities.

Suppose that the outcomes of coin flips $s_1, s_2, s_3, \ldots$ depend on each other such that s_2 tends to be equal to s_1, then s_3 tends to be equal to s_2, and so on. The question is how far the information from s_1 can travel in this one-dimensional system.

If the coin flips do not depend on each other, i.e., each s_i reproduces s_{i-1} with a 50% chance, then the outcome of a coin flip provides no information about the outcome of any other coin flip. Hence, information about s_1 is not transferred and the IDL is zero. See Figure S1a.

Suppose now that each outcome s_i is equal to s_{i-1} with a 75% probability. The second coin flip s_2 can infer the probability distribution of the outcome s_1 by using Bayes' theorem $p(s_{i-1} | s_i) \cdot p(s_i) = p(s_i | s_{i-1}) \cdot p(s_{i-1})$, which in this example means simply that s_1 is distributed 75%-25% over its two possible states. Using Eq. (1) we find that s_2 stores 0.19 bits about s_1, or in other words, 19% of the state s_2 is actually a reflection of the state s_1. The remaining 81% of its state is still randomness or noise, as before. Similarly, s_3 can use Bayes' theorem to find that its state is equal to s_1 with probability $0.75^2 + 0.25^2 = 68\%$, so according to Eq. (1) it received 0.046 bits of information from s_1. Clearly, the 1 bit of information about s_1 is imperfectly transferred through the system and eventually vanishes. See Figure S1b.

In the more general case there can be prior information about s_1 already stored in the system. This information is not transferred through dynamics but can be due to an external force or prior knowledge. For instance, let s_1 be the outcome of an unfair coin flip that is distributed 75%-25%. Even in the absence of interactions, each subsequent outcome s_i can already infer 0.19 bits of the state s_1. Information received due to interactions will be additional to this 'baseline' information, see Figure S1c.

We can now calculate the IDL of this example. All coins and their interactions are equivalent, so we expect a constant rate f of losing information at each subsequent coin flip:

$$f = \frac{I(s_1 | s_n)}{I(s_1 | s_{n+1})} = \frac{H(s_1) - H(s_1 | s_n)}{H(s_1) - H(s_1 | s_{n+1})} = \frac{(\frac{1}{2} - \frac{1}{2^n}) \log(\frac{1}{2} - \frac{1}{2^n}) + \frac{1}{2^{n+1}} \log \frac{1}{2^{n+1}} - \log 2}{(\frac{1}{2} - \frac{1}{2^{n+1}}) \log(\frac{1}{2} - \frac{1}{2^{n+1}}) + (\frac{1}{2} + \frac{1}{2^{n+1}}) \log(\frac{1}{2} + \frac{1}{2^{n+1}}) - \log 2}.$$

This rate is constant except for a small deviation for the lowest n, which we find by taking the limit of n using L'Hopital's rule:

$$\lim_{n \to \infty} f = 4.$$

In words, each subsequent coin flip s_i stores one quarter of the information that its predecessor s_{i-1} stores. We define the information dissipation length as the characteristic halftime of information, so in this case

$$\text{IDL of coin flips} = \frac{\log \frac{1}{2}}{\log \frac{1}{f}} = \frac{1}{2} \quad (2)$$

Following this intuition we measure the IDL in the one-dimensional system $r_1^t, r_2^t, \ldots$, where r_i^t is the return (relative difference) of a IRS of the ith maturity (1 year, 2 years, etc.) at time t. We assume that all interactions among subsequent rates are equivalent and that each rate stores information in a similar manner. In this case there is a constant rate f^t of losing information about r_1 at each time t, which leads to the IDL at time t through Eq. (2).

S2 How IDL differs from other leading indicators

The primary difference between IDL and other the alternative leading indicators analyzed here is that it filters out external correlation at each time point. The IRS prices across maturities are financial indicators which not only correlate amongst themselves, but may also correlate with external indicators such as the house-price index (HPI). A standard correlation computed between IRS prices therefore consists of two parts: their interdependence (cause-and-effect) and the external correlation they have in common. This is the case in the previously introduced spatial leading indicators[3–7].

As the causal relation between two IRS prices becomes more indirect (larger difference in maturity), the interdependence vanishes to zero, and the external correlation that they all have in common remains. IDL estimates the magnitude of this common external correlation at each time point, and computes the characteristic length scale of the decay of interdependence on top of it. This idea is illustrated in Figure S1c. It is analogous to the classical concept of correlation length, where the mutual information is used as correlation measure.

The second difference is the use of the mutual information measure to compute correlations between time series. Typically, Pearson-like and other linear correlation measures are used to construct leading indicators. In contrast, the mutual information measure is capable of detecting various non-linear forms of correlation as well[8].

S4 Robustness of IDL as a leading indicator

Estimating the IDL requires two parameters: the size of the sliding window w and the number of bins h used to construct the contingency table. At each time point we use the preceding w price values in two time series $X = x_1,...,x_w$ and $Y = y_1,...,y_w$ in order to estimate their mutual information $I(X | Y)$. The number of bins determines the price equivalence relation, i.e., which price values are considered equal in each sliding window. This is necessary to calculate the mutual information using the discrete version of Eq. (1), i.e., to estimate the joint probability distribution of the two price values using a finite set of observations. The values of each time series X and Y are first binned into h bins such that each bin contains approximately the same number of samples and then the contingency table is constructed, known as adaptive (product) partitioning[9].

The higher the sliding window size w, the more accurate can mutual information be estimated but the less sensitive it is to detecting short-term events or sudden changes. Therefore, w should be as low as permitted by the accuracy of calculating the mutual information.

If the number of bins h is too high then no price values will be considered equal, which means that each observed pair of prices is unique and the mutual information is invariably maximum. Decreasing the number of bins implies a lower sensitivity to small correlations of price fluctuations, so the number of bins determines the magnitude of price changes that are correlated. Here too is a trade-off between accuracy and sensitivity. We use a data-driven statistical procedure to select the order of magnitude of the h parameter in Section S4.1 within which we vary its value to study the robustness.

We show the IDL for a wide range of parameter values for w and h in Figure S2 through Figure S4. We show at least one value (50) for w which is 'too small' in order to indicate its lower limit. This shows in what range the window size should be chosen; we chose $w = 300$ as a tradeoff between noise and sensitivity for Figure 1 in the main text.

We find a remarkable robustness of the IDL indicator if w is chosen to be at least 200, and h is chosen to be less than 50. The shapes of the IDL curves differ only in details, and the timing of their peak is consistently at the Lehman Brothers bankruptcy.

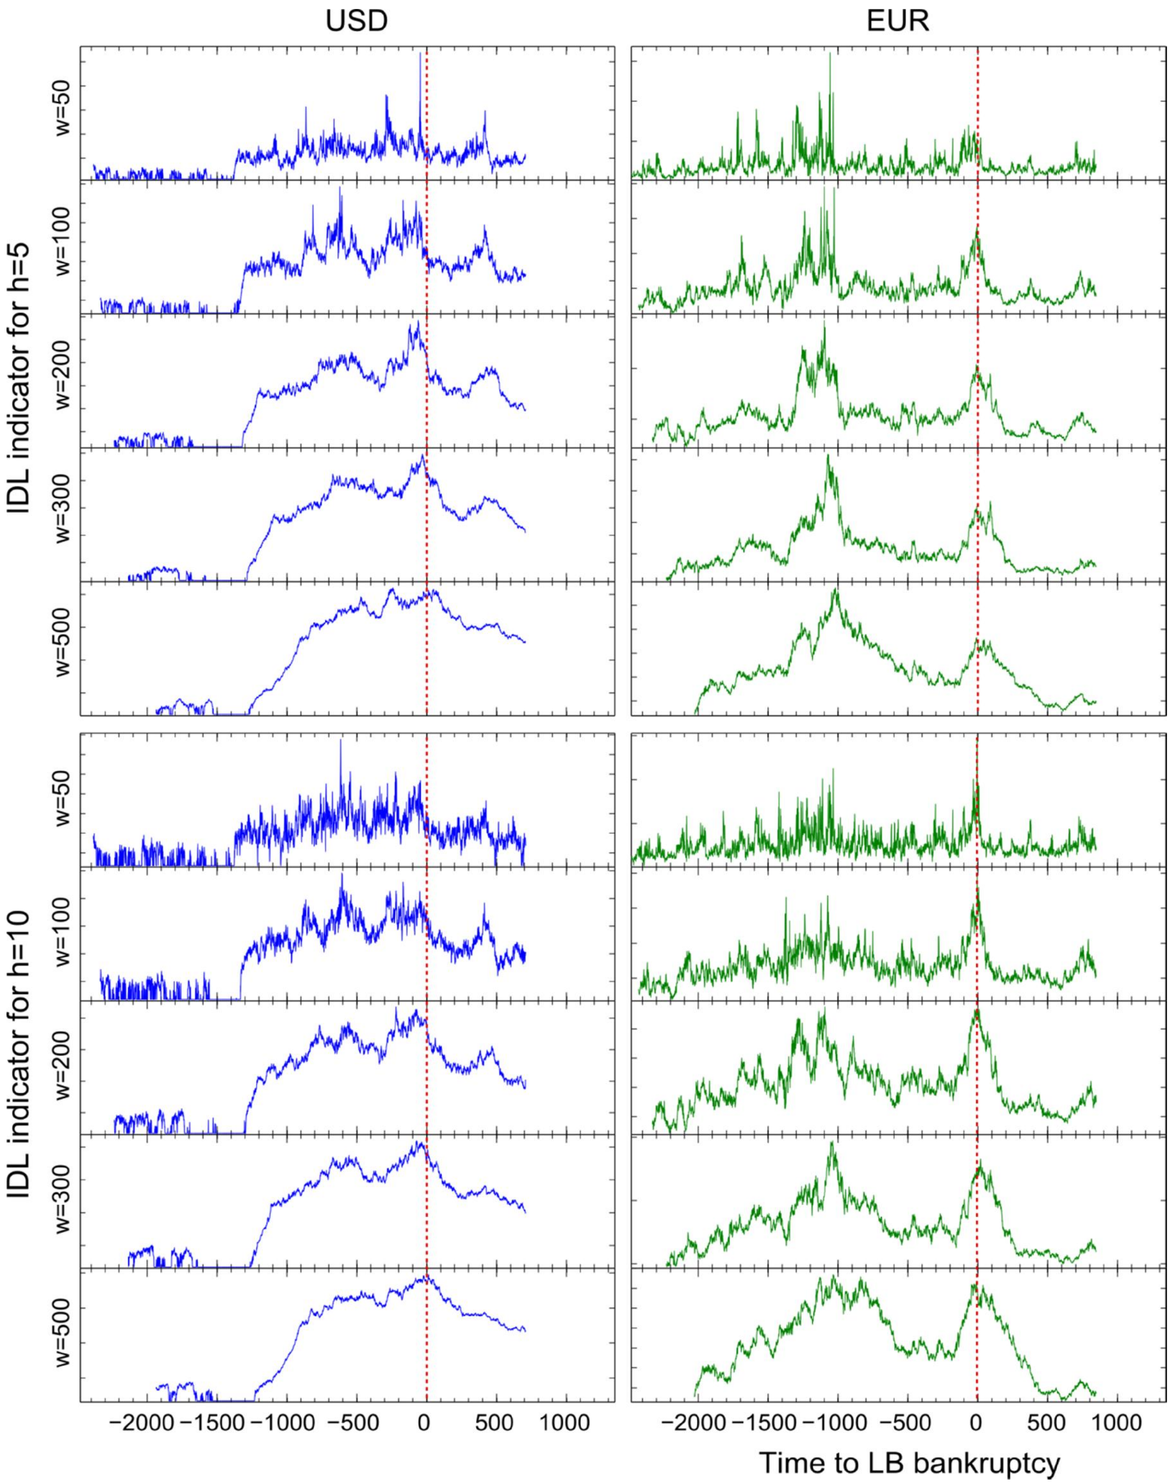

Figure S2: The IDL indicator for IRS rates in the USD and EUR currency, where 5 and 10 bins are used for the contingency table and sliding window sizes 50, 100, 200, 300, 500. Time point zero corresponds to the day of the Lehman Brothers bankruptcy.

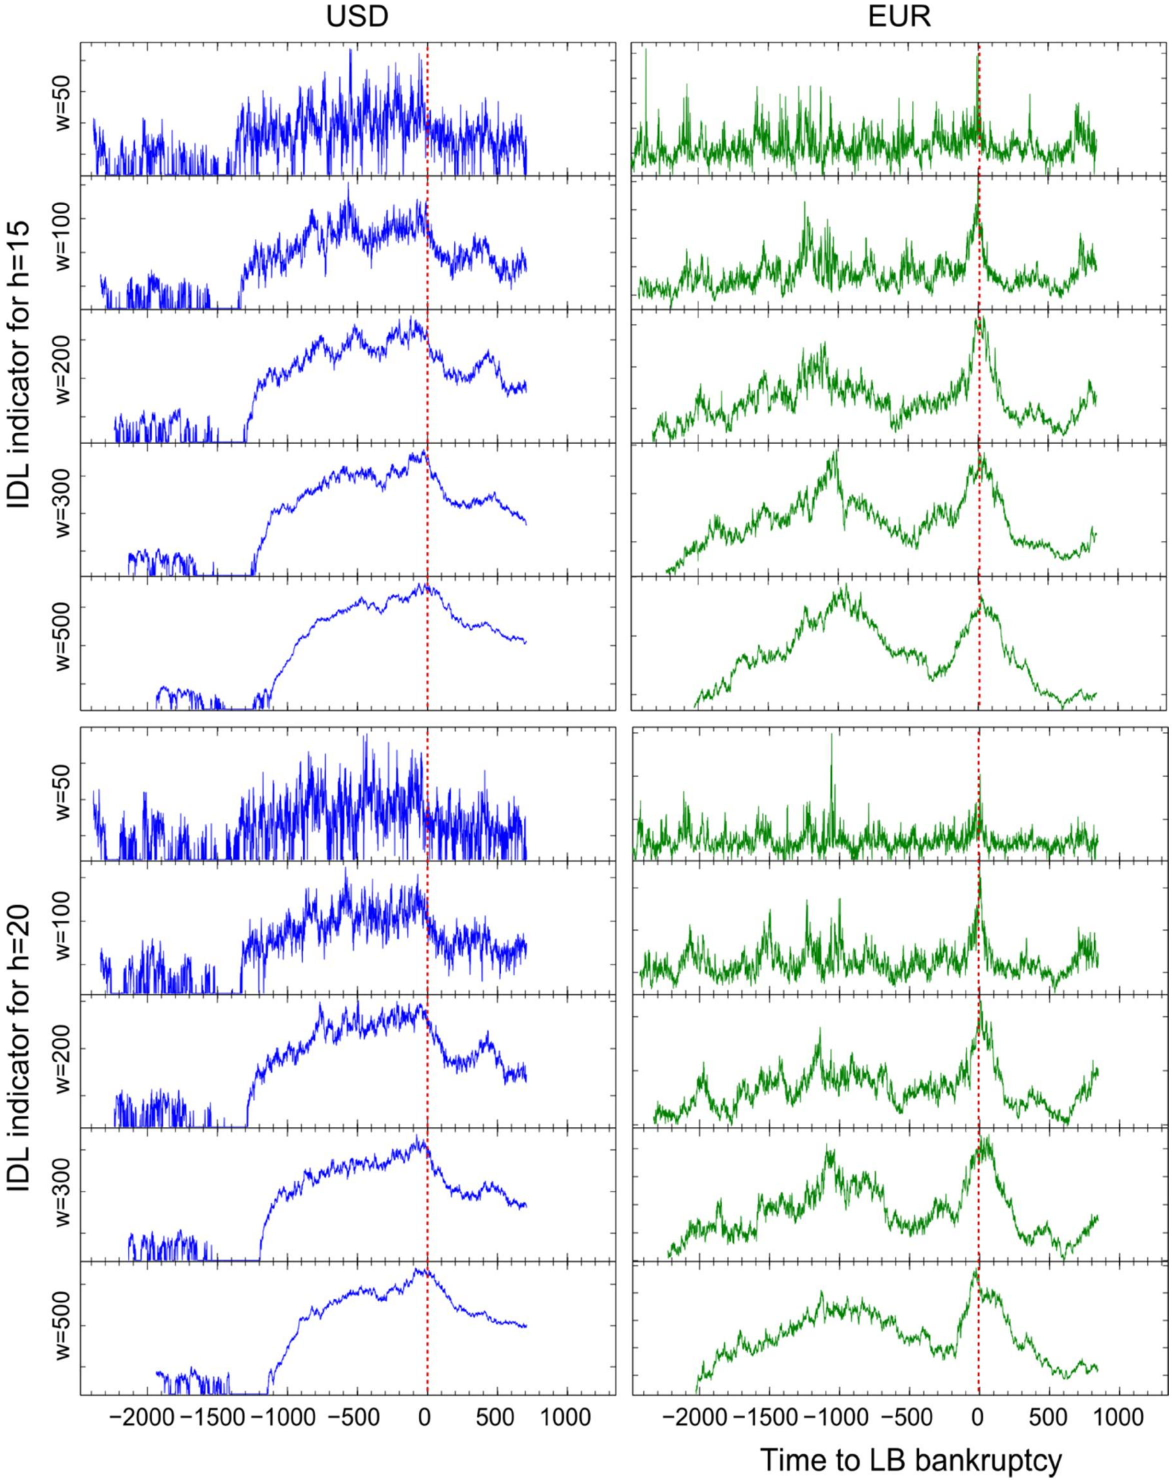

Figure S3: The IDL indicator for IRS rates in the USD and EUR currency, where 15 and 20 bins are used for the contingency table and sliding window sizes 50, 100, 200, 300, 500. Time point zero corresponds to the day of the Lehman Brothers bankruptcy.

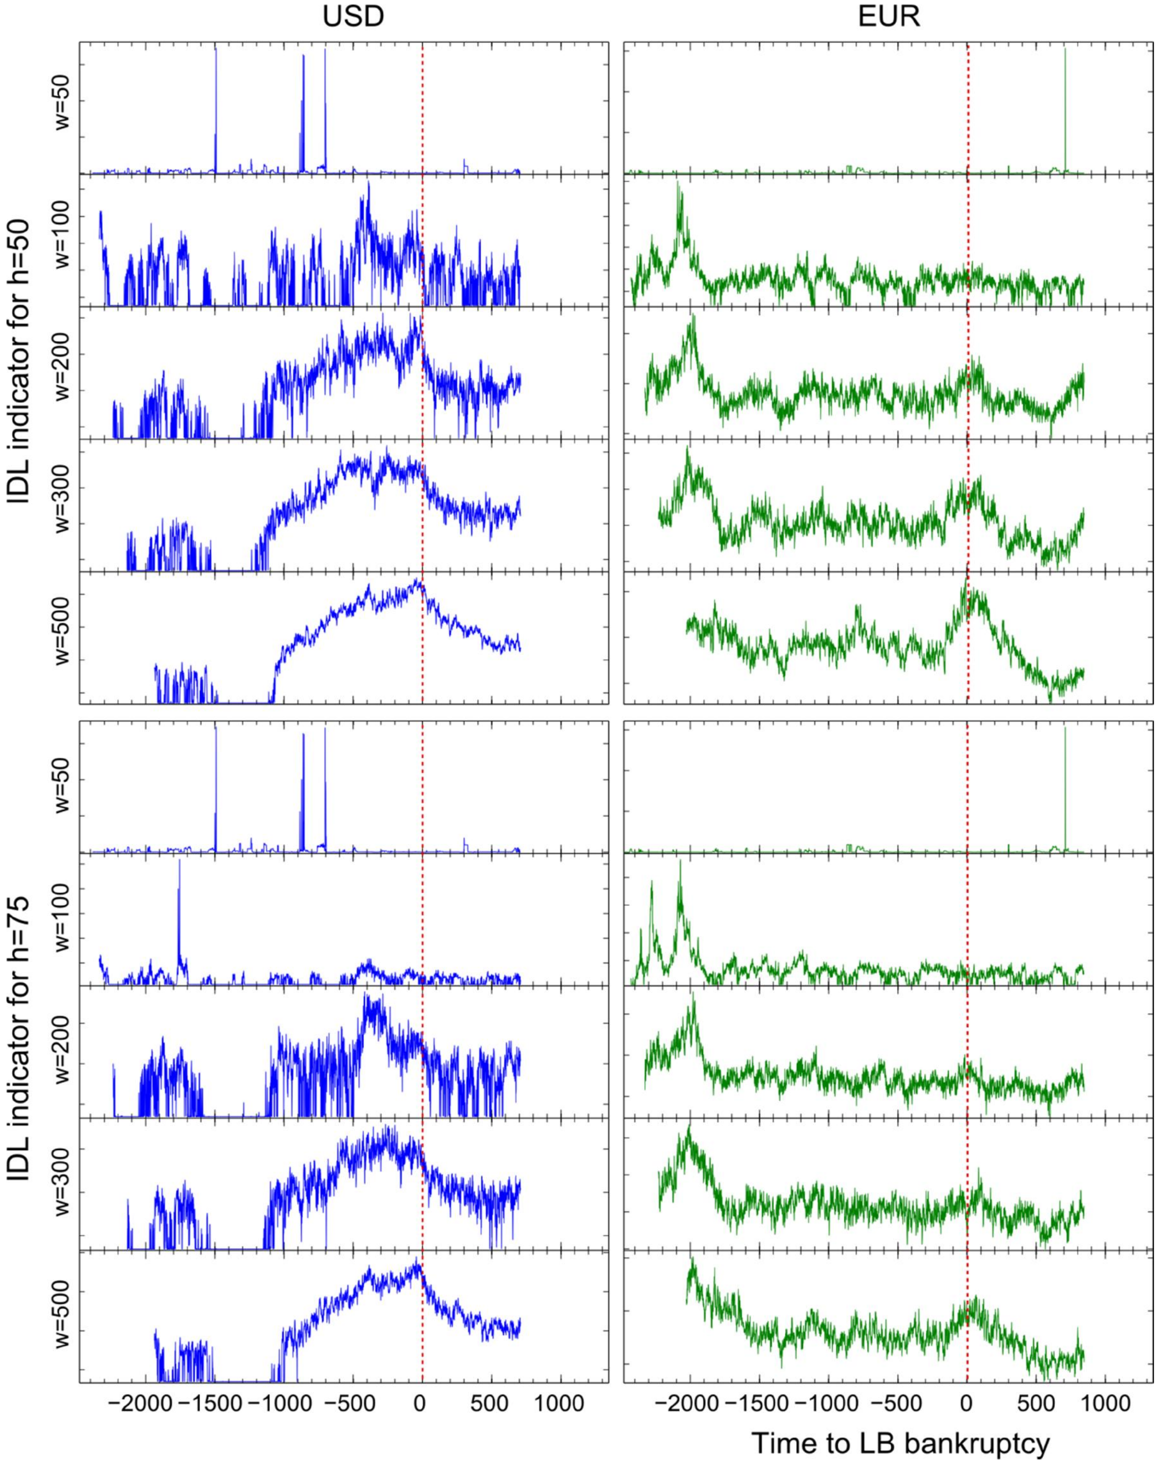

Figure S4: The IDL indicator for IRS rates in the USD and EUR currency, where 50 and 75 bins are used for the contingency table and sliding window sizes 50, 100, 200, 300, 500. Time point zero corresponds to the day of the Lehman Brothers bankruptcy.

S4.1 Choosing the number of bins based on the data

We follow a systematic procedure in order to find an optimal number of bins h based on the 'minimum description length' principle[10]. The parameter h induces a histogram estimator of the underlying (unknown) probability density of the set of w IRS returns at a particular time. Naturally, a good estimator follows the data closely while still being parsimonious. Intuitively, the estimator follows the data closely if it assigns high probability to dense clusters of values, and low probability to value ranges that contain only few samples. It is parsimonious if it is not too complex to describe, i.e., requires only a small number of characters to write down the data together with the estimator.

Rissanen[10] defined the formula $F(h)$ to calculate this description length, or 'stochastic complexity', as a function of the parameter h, as

$$F(h) = \sum_{i=1}^{h} n_i \log R_i + \log \frac{n!}{\prod_i n_i!} + \log \frac{(n+h-1)!}{n!(h-1)!} \quad . \tag{3}$$

Here, n_i is the number of data points that fall in the ith bin and R_i is the width of bin i. The value for h that minimizes this stochastic complexity is then considered the optimal value that should be used to summarize the last w data points at time t. We use $w = 300$ as determined in the previous Section and repeat this procedure for all windows $r_1^{(t-299)}, \ldots, r_1^{(t)}$ in the 1-year maturity USD data, where t runs through the time series with step size 10. The result is then a distribution of optimal values for h for our data, which serves as the range of values to be used for our IDL calculations in the previous Section.

Instead of computing $F(h)$ for the case of adaptive partitioning (variable bin width), which is the algorithm used to calculate the IDL curves, we compute $F(h)$ for the equidistant binning procedure, where each bin has the same width. The reason is that the stochastic complexity formula appears to handle duplicate values poorly, which for our data is an average of 10% of a randomly chosen window of 300 values. The consequence of duplicate values is a deviation from the uniform distribution of data points into bins, i.e., some bins will have more samples than average, and some other bins will have fewer samples. This makes the second term, which is code length required to encode the bin number for each data point, (much) smaller than it should be, such that its upward trend never outweighs the downward trend of the first term, which is the code length needed to place a data point within a bin's width given that it is known to which bin it belongs. The result is that for all windows we find an optimal h that is close to w or even exceeds it, which is clearly not a parsimonious description of the data so we deem this method incorrect. Using fixed width bins does not have this problem so we use this procedure as an approximation. We will compare the result with two additional rules of thumb from the literature in order to verify that this approximation is indeed suitable.

A typical curve of stochastic complexity $F(h)$ as function of h is shown in Figure S5. We estimate the location of the minimum by first smoothing the curve using a Gaussian kernel with a standard deviation of 5 data points. Then we find the global minimum value of this smoothed curve, identifying the optimal h for the particular window. The histogram of optimal values h_{opt} for all windows is shown in Figure S6.

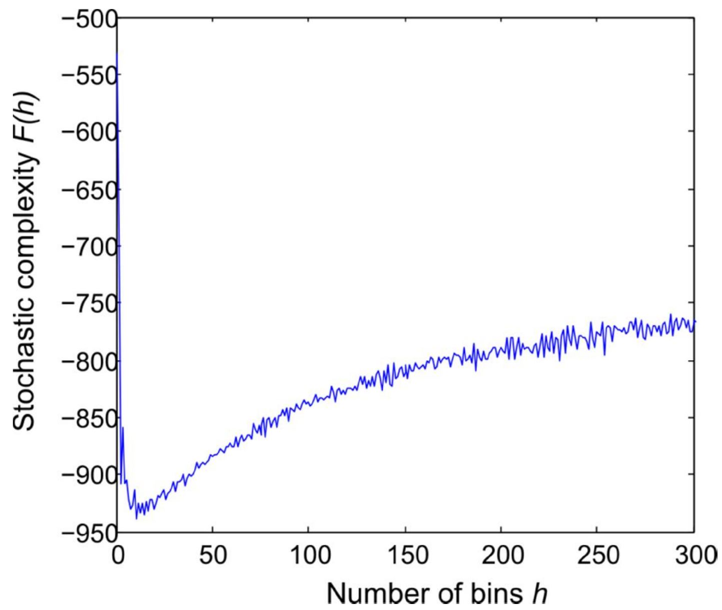

Figure S5: A typical curve for the stochastic complexity $F(h)$ as function of h for the equidistant binning procedure, using a sliding window size of $w = 300$. This curve is calculated from the last 300 returns available at the 1000th trade day in the 1-year maturity USD data. The optimal number of bins to use for calculating the mutual information at this time is identified by the minimum of the curve. The values on the y-axis have no absolute meaning.

Although for most windows we found an intermediate optimal h, for some windows $F(h)$ was monotonically decreasing and therefore yielded a very high optimal value. The overall average optimal h is 50.16 bins. Ignoring the incidental high values yields an average optimum of 17.26 bins. The standard deviation of the lower values is 1.58. Recall that this estimate only approximates the magnitude of the optimal value since we use the equidistant binning procedure instead of the equiprobable binning procedure.

These results are in line with three independent rules of thumb for choosing the optimal number of equidistant bins. Mosteller and Tukey[11] suggest to use $n^{1/2}$ as the number of bins, which in our case is 17.32. Another rule of thumb is due to Bendat and Piersol[12] who suggest to use $1.87(n-1)^{0.4}$ which in our case becomes 46.86. Finally, Terrell and Scott[13] suggest to use $(2n)^{1/3}$ bins, which becomes 8.43. In conclusion we select the range of 5 through 75 for the number of bins in our study, and show the IDL curves for $h \in \{5,10,15,20,50,75\}$ in the previous Section. Indeed we observe a tendency for the IDL signal to deteriorate as h grows to 50 and 75, confirming that we selected a suitable range.

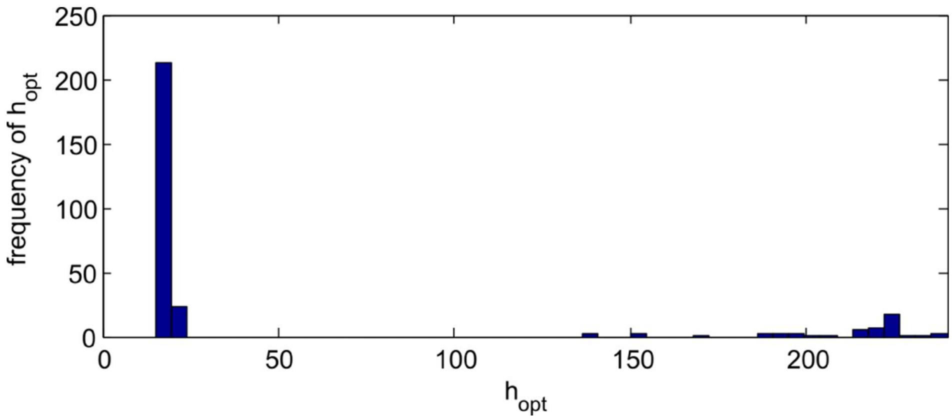

Figure S6: The histogram of optimal values for h for the entire 1-year maturity USD data (3145 returns), estimated from 300-day overlapping windows placed at 10-day intervals. The high values are close to the number of unique data points in the corresponding windows, which we interpret as outliers of the procedure.

S5 Comparison with previously introduced leading indicators

S5.1 Critical slowing down in IRS rates

The effect of critical slowing down[14] can be measured by the coefficient of a first-order autoregression of the fluctuations of a signal[3–5,15–18]. Calculating this coefficient requires two parameters: the size of the smoothing kernel, which de-trends the signal, and the size of the sliding window, which is used to compute the autoregression. Here we investigate whether there is a set of parameter values for which the critical slowing down can provide a clear leading indicator of the Lehman Brothers bankruptcy. In Figure 2 in the main text we show a representative set of results, where the smoothing kernel has a standard deviation of 5 trade days and the sliding window to compute the autoregression was 500 and 1000 trade days. Here we show the results for a wide range of parameter values.

The smoothing kernel is used to filter long-term price trends from the time series. We use a Gaussian smoothing kernel, following e.g. Dakos et al.[15]. The smoothing kernel is used to remove the long-term trend from a signal, because the effect of critical slowing down is detected in the short-term fluctuations of the time series: it is the time it takes for the price value to return to its long-term trend after a small perturbation. In effect we compute a running weighted average of each time series, where each price value becomes the weighted average of its neighbors, and subtract it from the original time series to obtain the de-trended signal. The weights are Gaussian distributed and the width of the distribution is the free parameter. Figure 2 in the main text was created using a Gaussian kernel with a standard deviation of 5 trade days, as in Dakos et al.[15]. Here we show the first-order autoregression coefficient for the parameter values 5 and 10.

At each time point the autoregression coefficient is calculated using the preceding w price values, where w is the size of the sliding window. The higher the value of w, the more accurate can the coefficient be calculated but the less sensitive it is to short-term effects. The first drawback that we find is that the calculation of the coefficient requires a considerably larger sliding window than for calculating the IDL. Where the IDL indicator starts to be meaningful at a size of about 100 trade days, the autoregression coefficient requires a window size that exceeds 500 trade days. This problem has already been recognized by others[3–5,16]. Figure 2 in the main text was created using a sliding window of 1000 trade days and a Gaussian smoothing kernel with a standard deviation of 5 trade days. In Figure S7, Figure S8, and Figure S9 we show the first-order autoregression coefficient of the de-trended IRS rates for the sliding window sizes 500, 1000, 1250, and 1500 trade days, a Gaussian smoothing kernel with standard deviations 5 and 10 trade days, for a representative sample of maturities: the first, second, fifth, and tenth maturity of each dataset. For the USD data this corresponds to maturities of 1 year, 2 years, 5 years, and 10 years; in the EUR data to 2 years, 3 years, 6 years, and 12 years.

We find a wide diversity of behaviors of the indicator near the Lehman Brothers bankruptcy, including upward trends (‘critical slowing down’), downward trends, and sudden increases and decreases shortly after the bankruptcy. We find that an early warning could be issued only in the EUR market for the 2-year maturity IRS, for a sliding window of around 1000 trade days but no more than about 1250 trade days. Our conclusion is that this indicator is not a reliable early warning signal.

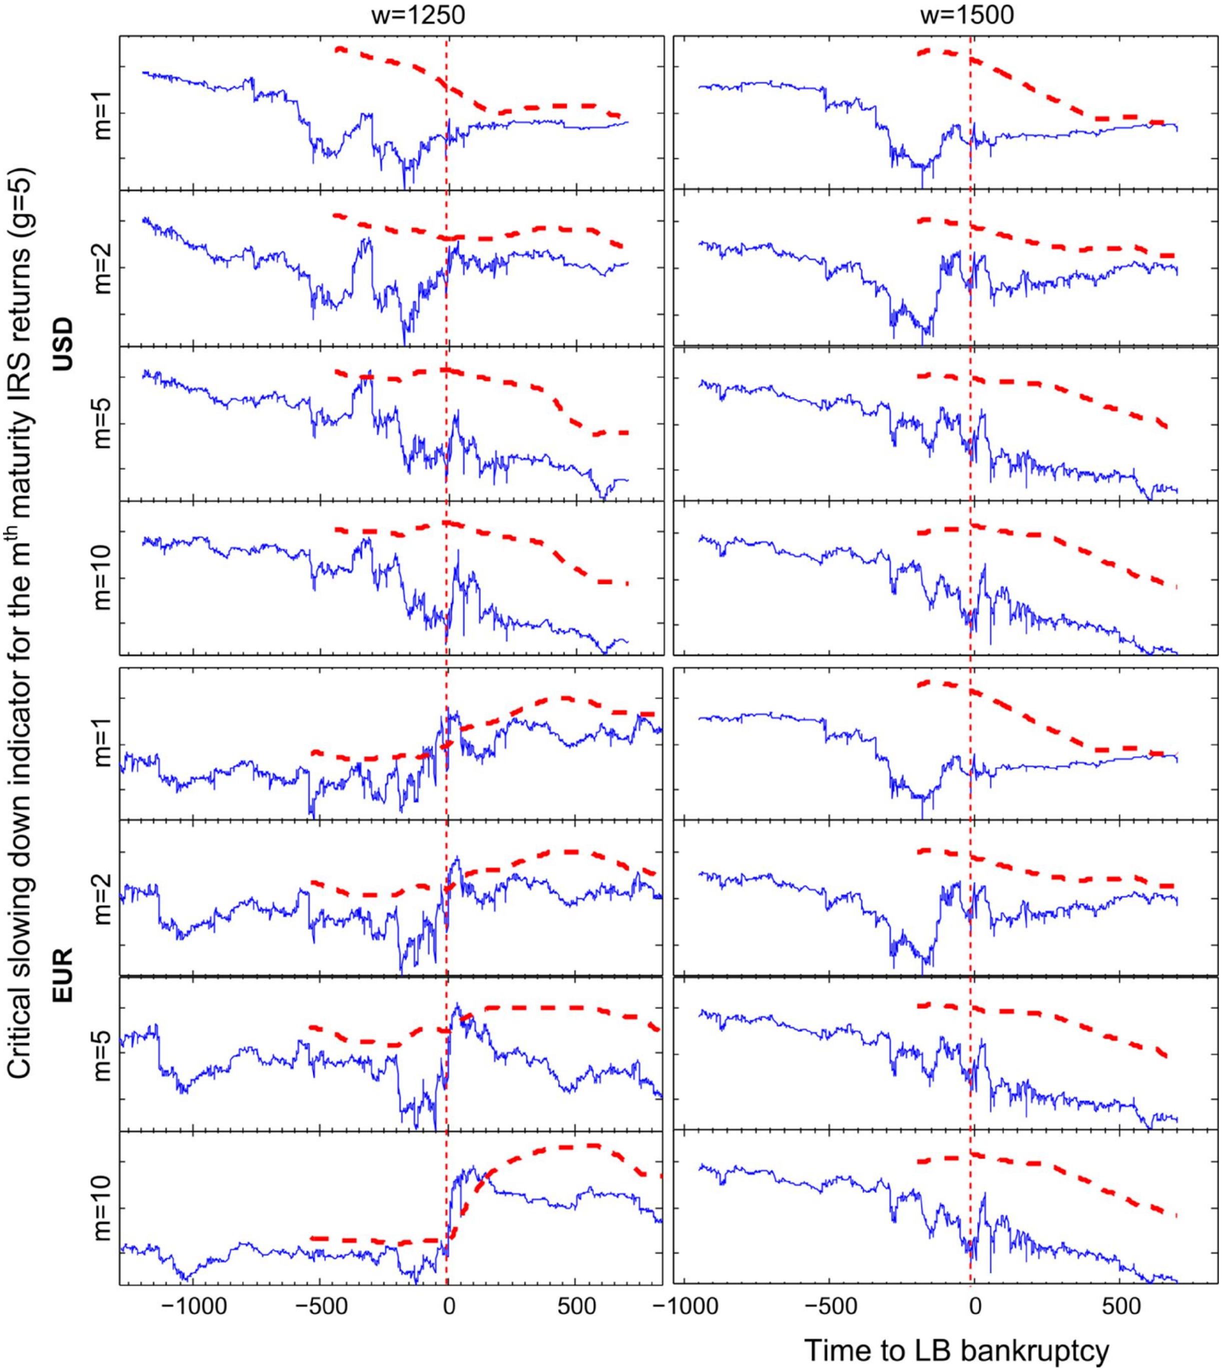

Figure S7: The first-order autoregression coefficient of the fluctuations of the IRS levels for the first, second, fifth, and tenth maturity in the USD and EUR markets. The red dashed line is the warning threshold, calculated as two times the standard deviation above the mean of 750 trade days, which is the same as for the IDL indicator. This indicator of critical slowing down is calculated for a sliding window size of 1250 trade days (left) and 1500 trade days (right). The Gaussian smoothing kernel used had a standard deviation of 5 trade days.

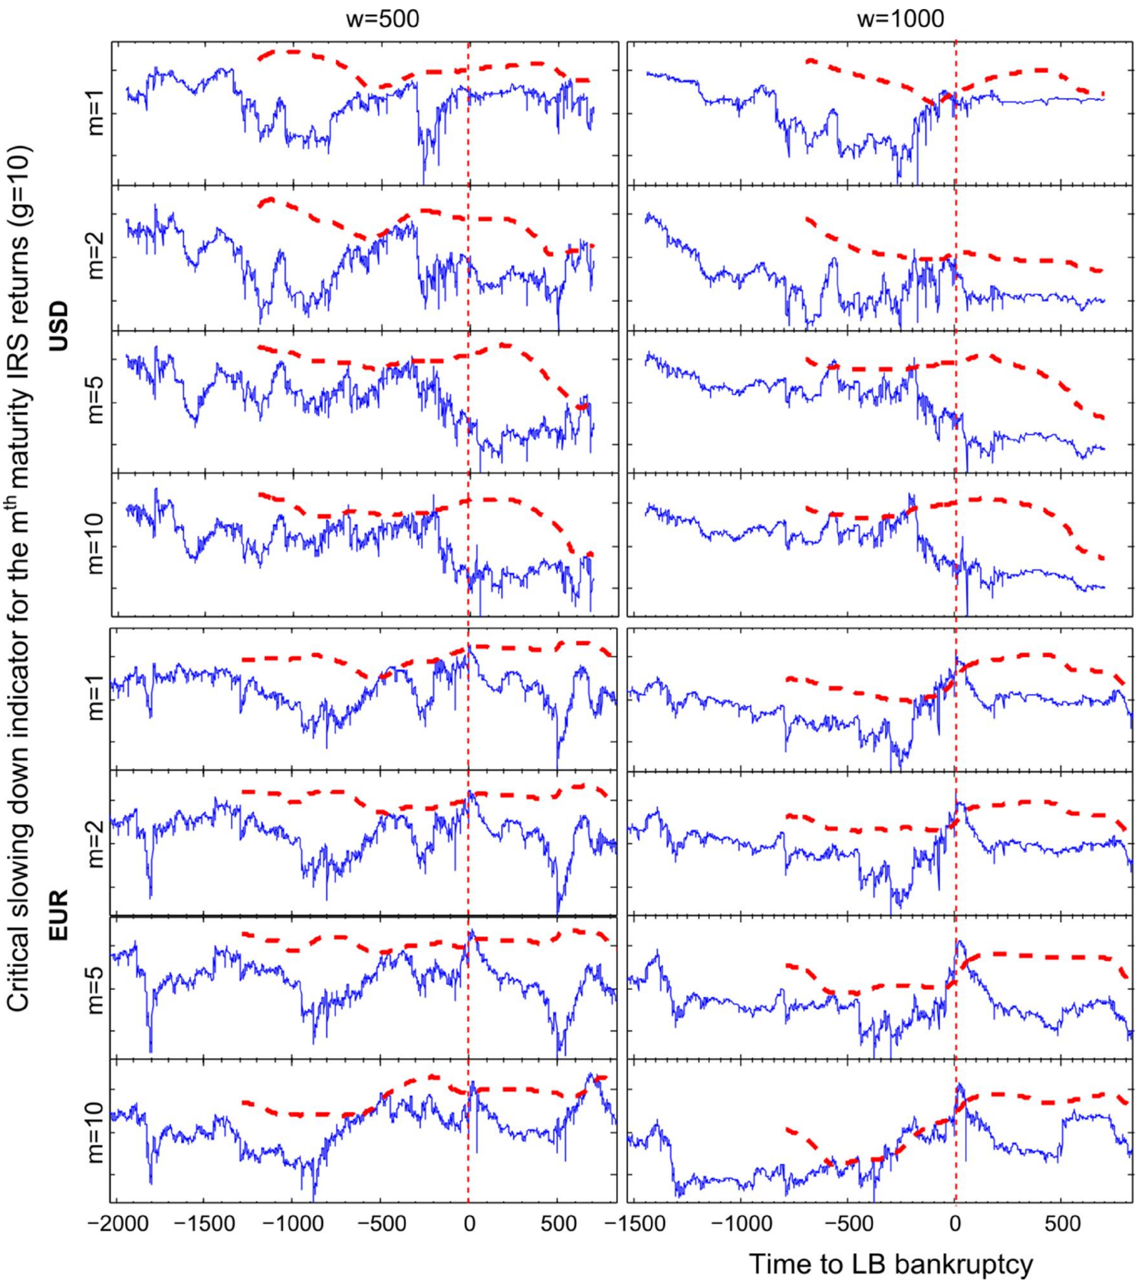

Figure S8: The first-order autoregression coefficient of the fluctuations of the IRS levels for the first, second, fifth, and tenth maturity in the USD and EUR markets. The red dashed line is the warning threshold, calculated as two times the standard deviation above the mean of 750 trade days, which is the same as for the IDL indicator. This indicator of critical slowing down is calculated for a sliding window size of 500 trade days (left) and 1000 trade days (right). The Gaussian smoothing kernel used had a standard deviation of 10 trade days.

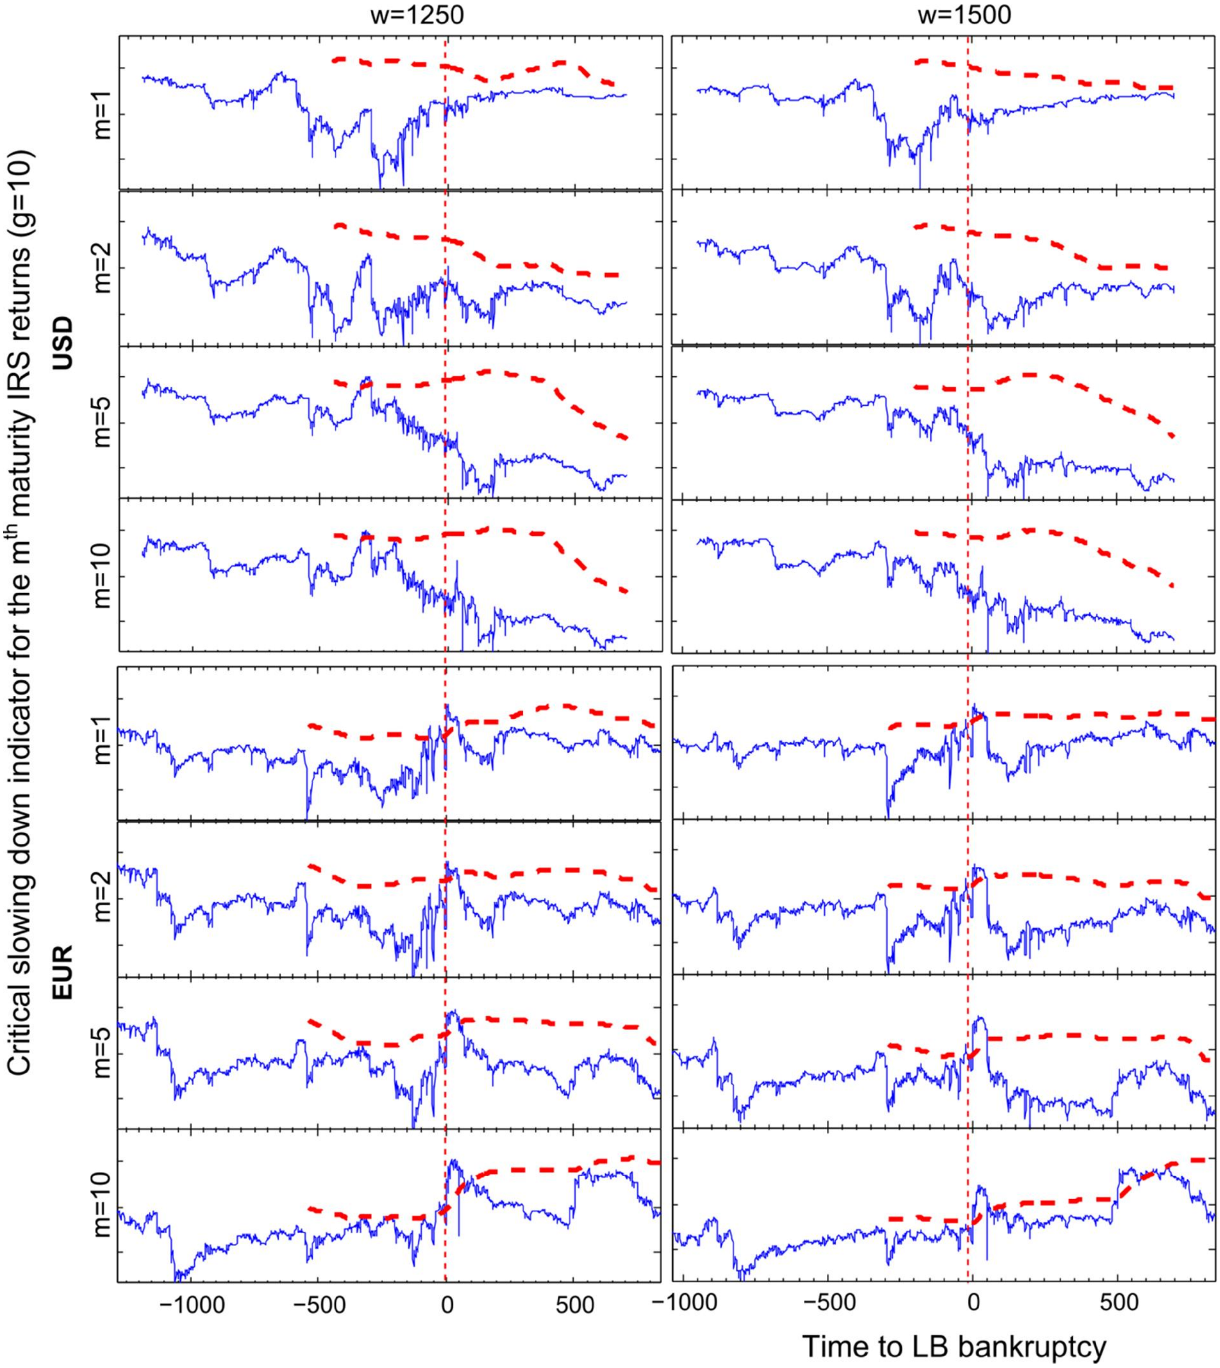

Figure S9: The first-order autoregression coefficient of the fluctuations of the IRS levels for the first, second, fifth, and tenth maturity in the USD and EUR markets. The red dashed line is the warning threshold, calculated as two times the standard deviation above the mean of 750 trade days, which is the same as for the IDL indicator. This indicator of critical slowing down is calculated for a sliding window size of 1250 trade days (left) and 1500 trade days (right). The Gaussian smoothing kernel used had a standard deviation of 10 trade days.

S5.2 The onset of a LIBOR-OIS spread as leading indicator

During the build-up of the recent crisis in 2007 a LIBOR-OIS spread, the so-called 'basis' in swap contracts emerged. The prices of the same swap but with different frequencies of variable payments had always been roughly equal, but around August 2007 the prices of swaps with less frequent payments started to increase. This is a very significant event that had never occurred before. The phenomenon is a symptom of calculating the risk of default of a financial institute, or in other words, a lack of trust in the stability of financial institutes[19]. The price difference is essentially an insurance premium to compensate the risk that the variable interest payer would default during the swap contract's duration. Such a default had hardly been considered before 2007.

The question is whether the onset of a basis in interest rates can be used to anticipate the bankruptcy of Lehman Brothers. Therefore we interpret the basis as a leading indicator and test whether it could be used as an early warning for the bankruptcy. As warning threshold we use the same definition as for all other leading indicators, namely three times the standard deviation above the mean of a sliding window of s=400 trade days. In addition, we test the case for a sliding window of s=200 and s=800 trade days. As a basis we use the daily differences between the 3-months (3M) LIBOR rates, which is based on a financial contract with one single payment, and the 3-months overnight indexed swaps (OIS), which have daily payments. The first reason for taking the difference with the 3M swaps is that it is the most liquidly traded; the second reason is that it is the same frequency as the USD data analyzed in the main text and here.

A 3M-OIS swap is a different type of swap from an IRS which party A and B can negotiate. Party A pays x USD to B immediately after signing the contract, and B pays back $x + \text{LIBOR} \cdot 3/12 \cdot x$ to A after three months. In such a swap there is no notional exchange; only one fixed rate is exchanged with one floating rate. This floating rate is the average daily interest rate cumulated over the period of three months. There is also more credit risk in such swaps: if B defaults then A loses its x USD.

See Figure S10 for the 3M-OIS interest rate and the three different warning thresholds. Although in each case there are two periods where a warning is issued, neither warning can be used to anticipate the Lehman Brothers bankruptcy. The first warning that lasts at least two days are at -291 trade days (s=200, lasting 33 days), -285 (s=400, lasting 29 days), and -285 (s=800, lasting 47 days). The basis may arguably be warning for a significant financial event around this time [19], however it is too early and too short-lasting to be used to anticipate the Lehman Brothers bankruptcy. The second warning is consistently too late (starting 3, 9, and 4 trade days after the bankruptcy), so it could be interpreted more as a consequence of the bankruptcy rather than anticipating it.

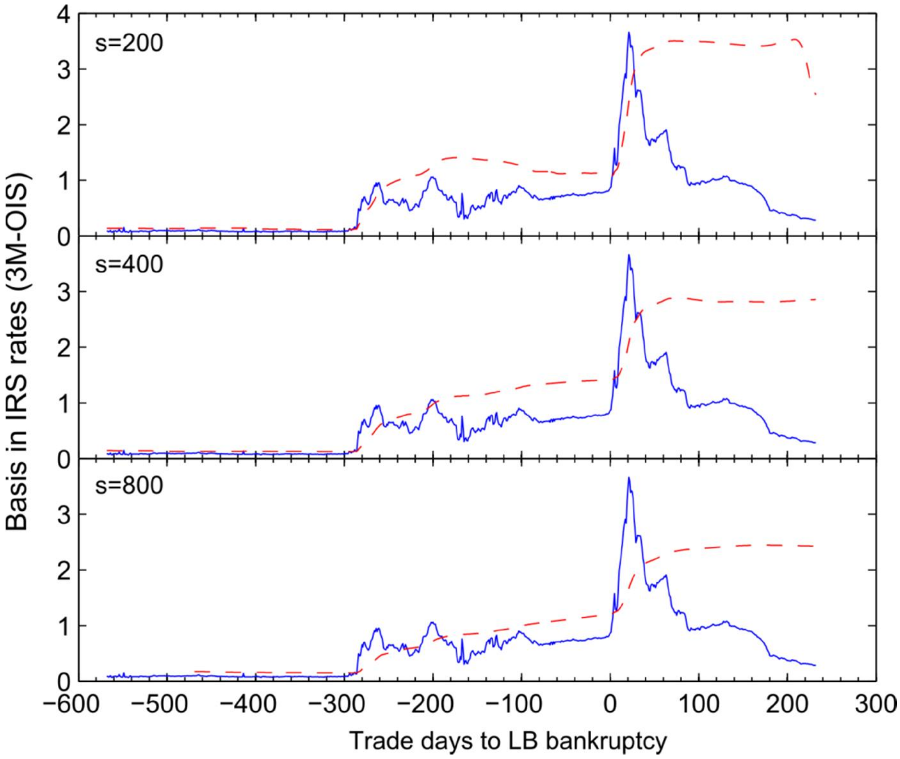

Figure S10: The basis in the USD IRS rates in the time span March 2006 through November 2009. The red dashed curve is the warning threshold, computed as three times the standard deviation above the mean of a sliding window of s=200, s=400, and s=800 trade days respectively. As a basis we use the daily differences between the swap rates with a variable payments frequency of three months (3M) and overnight indexed swaps (OIS), which have a daily frequency. Time 0 on the x-axis corresponds to the day of the Lehman Brothers bankruptcy.

S5.3 Critical slowing down in IRS spread levels

A traditional financial indicator is the 'spread' of (in this case) IRSs across maturities[20]. It is already evident from the IRS prices plot in Figure 1 in the main text as well as the cross-maturity variance plots in Figure 3 that the spread levels themselves do not provide an early warning for the Lehman Brothers bankruptcy. Nonetheless, since they are often used as underlying indices in complex interest rate derivatives ('spread options' [21]), it is possible that the critical slowing down (CSD) indicator applied to the spread levels provides an early warning signal. We already showed in Section S5.1 that the CSD does not anticipate the bankruptcy when it is applied to the original IRS levels.

We compute the spread levels as the daily differences of IRSs of various maturities compared to IRSs with a 1-year maturity (USD) or 2-year maturity (EUR). In other words, in USD the smallest spread is the 2-year IRS rate minus the 1-year IRS rate, then the 3-year IRS rate minus the 1-year IRS rate, etc. Next we calculate the first-order autoregression coefficient as in Figure 2 in the main text with a sliding window of 1000 and 1500 trade days and a Gaussian smoothing kernel with a standard deviation of 5 trade days. The warning threshold is computed as two standard deviations above the mean of 750 trade days, as before. The results are summarized in Figure S11 for representative parameter values.

Interestingly, the effect of critical slowing down is more apparent in the spread levels than in the original IRS rates, which were shown in Section S5.1. Of the 16 panels we find an early warning in two panels, namely for the USD spreads for $w = 1500$, for the 2y-1y and 5y-1y spreads. However we also observe all types of other possible behaviors of the indicator, especially for different window sizes, but even for the same parameter values for different maturities (3y-1y and 10y-1y). More research into this indicator applied to the IRS data is needed to understand why it is so sensitive to the choice of parameters and maturity.

Lastly we note that the sliding window size that generates an early warning is 1500 trade days, which means that each regression is performed on data of about 6 calendar years. The indicator can therefore measure only long-term trends of stability in the financial markets and is not sensitive to short-term critical transitions such as the Lehman Brothers bankruptcy.

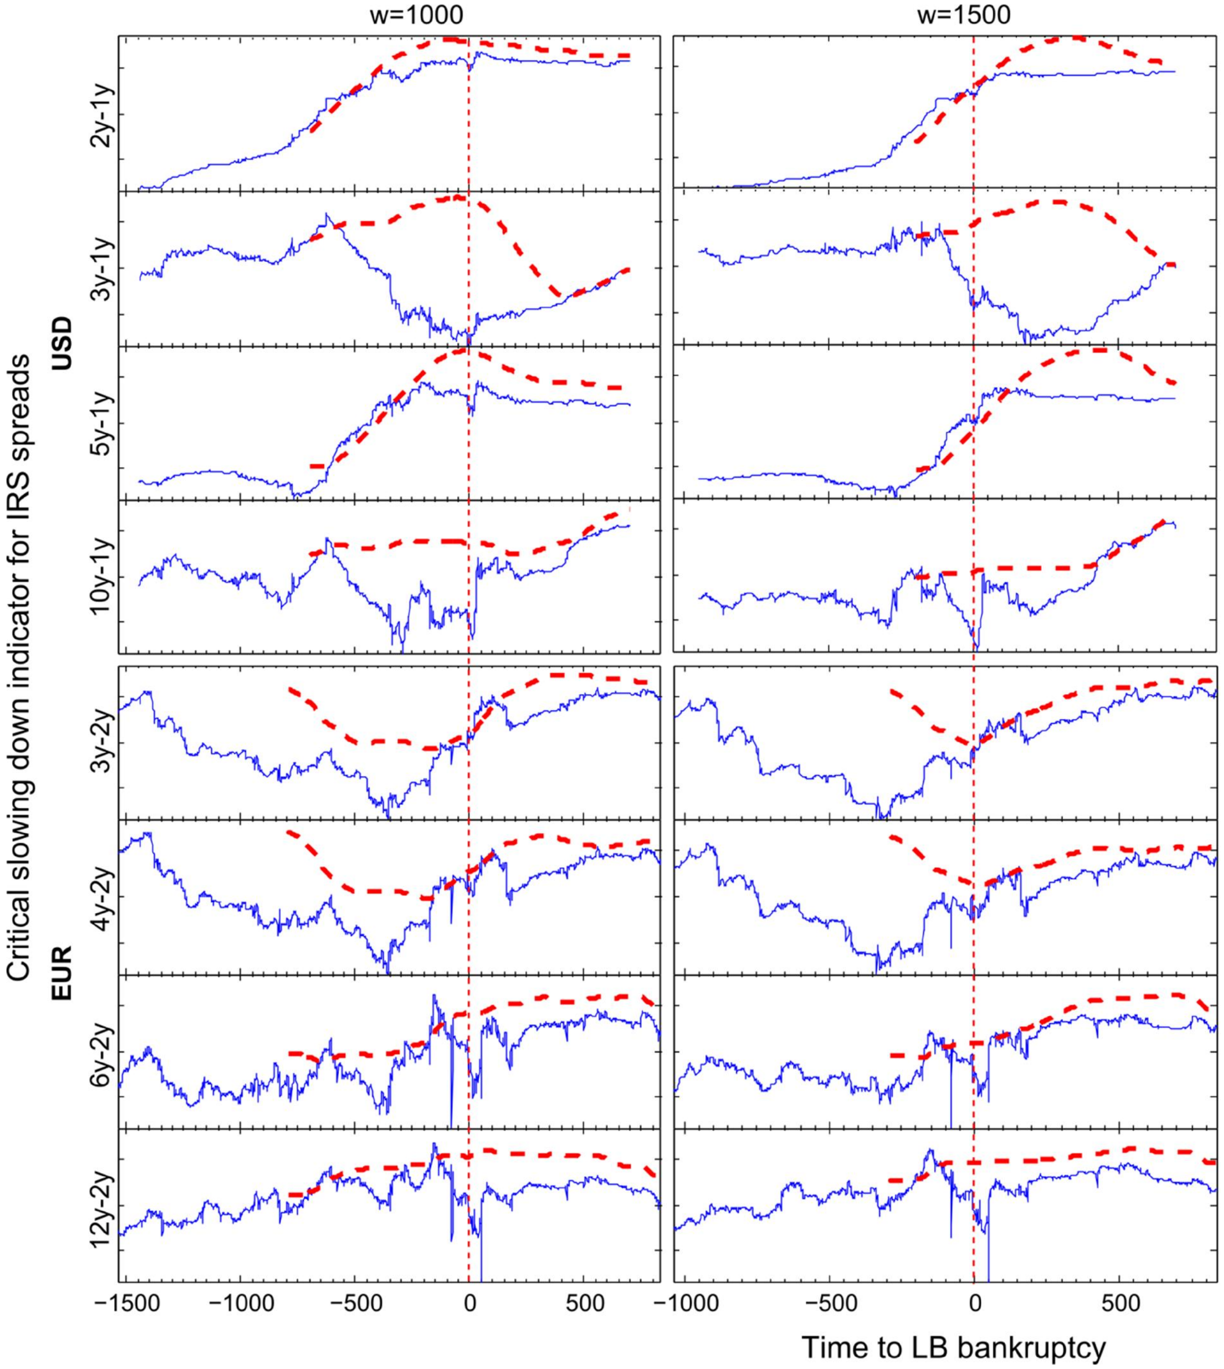

Figure S11: The first-order autoregression coefficient of de-trended IRS spread levels (sliding window sizes 1000 and 1500) and a warning threshold computed as two standard deviations above the mean of 750 trade days. A caption such as '2y-1y' denotes the spread level between a 2-year IRS and a 1-year IRS, i.e., the daily difference of their prices. Time 0 on the x-axis corresponds to the day of the Lehman Brothers bankruptcy.

S6 Verification of IDL using generated time series

Calculating the IDL of generated time series allows us to address two questions. Firstly, does the IDL curve contain peaks even if the time series are uncorrelated (false alarms)? Secondly, does the IDL curve contain a peak at the time where we let the time series correlate?

S6.1 Generating the time series

The generated data consists of 15 time series which consist of 3145 real-valued elements. These dimensions are equal to that of the USD IRS data so that the verification in this section is as comparable as possible to the real data used in the main text.

The first time series, i.e. the first 'maturity', is a copy of the 1-year USD IRS returns. Each subsequent i^{th} time series is a vector of random values, with a mean and standard deviation that equals that of the corresponding i^{th} maturity IRS data in the USD currency. The exception to this rule is a range of 500 elements starting at the 2000$^{\text{th}}$ element, where an artificial correlation is introduced as follows. The values of the elements 2000,…,2499 in the i^{th} time series, excluding $i = 1$, is a randomized copy of the values of the $(i-1)^{\text{th}}$ time series. The randomization is an added noise factor that has a zero mean and a standard deviation s, which is the free parameter.

S6.2 Results

The lower the value of s, the higher the correlation between subsequent time series, and expectedly the higher the IDL. In other words, the higher the value of s, the less information about the first time series propagates through subsequent time series. Therefore we expect that the IDL peaks during the 500 correlated 'trade days' with a magnitude that decays for increasing s. In the absence of correlations, we expect that the IDL curve contains no discernible peaks.

In Figure S12 we show the IDL of the generated data for $s = \{0.1, 0.5\}$ and for the case of uncorrelated time series. The IDL is calculated using a sliding window size of $w = 300$ and number of bins $h = 10$, matching the parameter used in the main text. The x-axis is equal to the USD IDL plot in Figure 1, and the period of serial correlation of 500 trade days is indicated by the grey area.

Indeed we observe that the IDL peaks during the period of correlated elements among subsequent time series, confirming our hypothesis that the IDL is capable of detecting correlations between subsequent maturities. The delay of the peak compared to the start of the grey area is expected because a sliding window is used, and no preceding gradual onset of correlation is generated.

Further we observe that the IDL does not contain significant peaks in the absence of correlations, suggesting that the IDL curve does not tend to generate 'false positive' warnings, i.e., warnings in the absence of correlated time series. We observe spurious IDL values not exceeding a value of 0.2, which we also observe in the initial low-IDL period in the USD market in Figure 1 in the main text. This suggests that these fluctuations in IDL in Figure 1 do not necessarily indicate small instabilities in the market during this period.

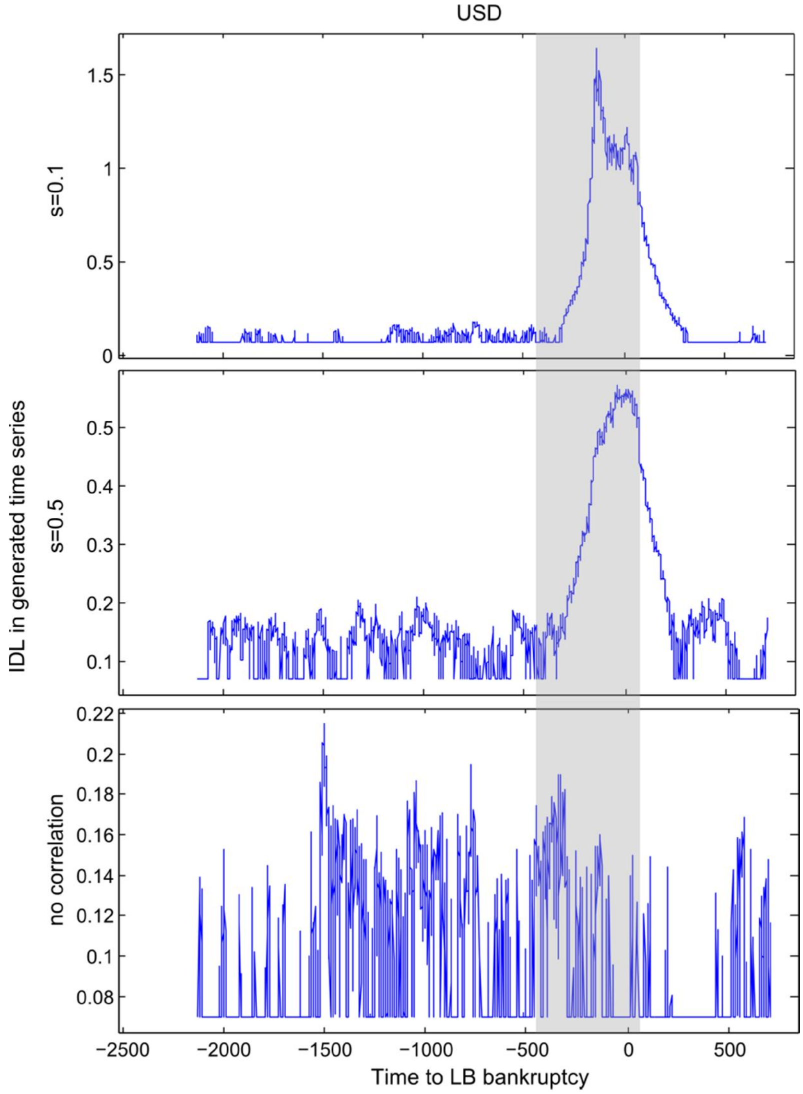

Figure S12: The IDL indicator computed of generated time series. The time series are uncorrelated except during the time points $-439,\ldots,39$ (grey area; the 2000th trade day until the 2500th); the higher s, the weaker the serial correlation. The sliding window size is $w=300$ and the bin sizes $h=10$, corresponding to the parameter values used in the main text.

S7 Choosing a threshold for warning signals

In order to issue warnings for imminent critical events in an objective manner a warning threshold must be defined. On a given trade day the warning signal is 'on' if the IDL indicator is higher than this threshold, and it is 'off' otherwise. Typically, the threshold is calculated as a number of standard deviations above the long-term mean of the leading indicator in order to distinguish structural increases from fluctuations. This definition requires two parameters: the size of the sliding window s used to compute the 'long-term' means and standard deviations, and the number of standard deviations n to add to the mean. A universal choice for these parameters does not exist. Therefore we fit n and s to a desired magnitude of market crashes that is detected.

The magnitude of the market crashes that a leading indicator detects has a positive relation with the mean time between warnings. That is, the longer the duration until the next warning, the less sensitive is the warning signal to small peaks in the leading indicator, and consequently the larger the market crash must be in order to let the leading indicator exceed the warning threshold. The reasoning behind this relation is the widespread evidence that market crashes are not statistical outliers: instead they are part of a continuous, fat-tailed size distribution of critical events[22,23], in which small events are abundant and large events are increasingly rare.

We estimate the mean 'time-to-warning' for the IDL indicator for a range of values for n and s by extrapolating the IDL indicator of Figure 1 in the main text. The IDL curves in this figure are too short to reliably estimate the mean time-to-warning it induces, so we generate longer time series with the same autoregressive coefficients as the IDL curves. In Figure S13 we show the autoregressive coefficients a_k for both IDL curves up to lag $p=10$, using the autoregressive model $\mathrm{AR}(p)$ defined by

$$\sum_{k=0}^{p} a_k y_{t-k} = \varepsilon_t \ ,$$

where ε_t is zero-mean white noise and the coefficients are normalized so that $a_0 = 1$.

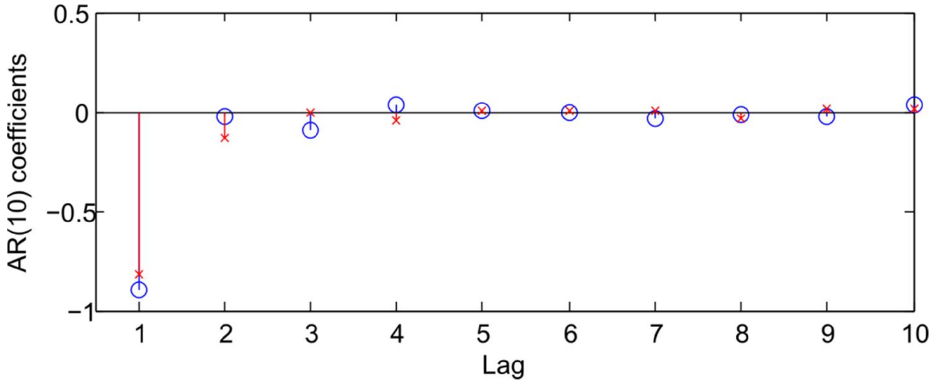

Figure S13: The autoregression coefficients of the IDL curves for the USD (red crosses) and EUR (blue circles) markets, up to lag 10. The IDL curves used are those in Figure 1 in the main text.

For the extrapolation we choose the coefficients of the IDL curve for the USD data. This is reasonable not only because the Lehman Brothers bankruptcy occurred in the U.S., but also because the choice of n and s seem more crucial for the USD market. In the EUR market the bankruptcy occurred during a period of relative stability, so the preceding onset of instability is distinct and easily detected. In the USD market, on the other hand, the bankruptcy is part of a larger period of growing instability, so whether or not the onset of 'instability-inside-instability' would be detected is more dependent on the parameter values. Based on Figure S13 we find that only first two coefficients are significantly different from zero, so we compute the coefficients for an $\mathrm{AR}(2)$ model, finding $a_1 = -0.8316$ and $a_2 = -0.1681$. We use this $\mathrm{AR}(2)$ model to generate an IDL curve for an equivalent of 500 years, assuming that each year consists of 250 trade days.

We find that warnings are highly clustered, see Figure S14. By far most warnings are less than a quarter year apart ($250/4 = 62.5$ trade days). Observing the two insets in Figure 1 in the main text, this observation makes intuitive sense: during a period of increasing instability, the leading indicator will reach the threshold and may be alternatingly above and below the threshold. This happens because the threshold grows along with the indicator.

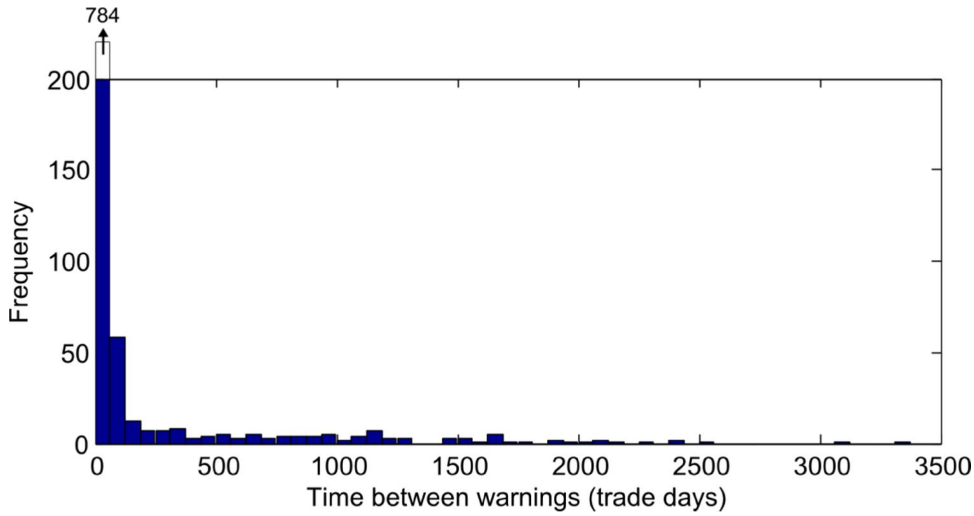

Figure S14: The distribution of the number of trade days between two warnings, generated by the simulated AR(2) IDL curve which spans 125000 trade days. The width of each bar is 250/4 = 62.5 trade days, which corresponds to a quarter year. The frequency of the time between warnings being less than 62.5 trade days is 784, which falls outside of the plotted region in order to make the rest visible. Within this left-most group, the number of warnings that are consecutive (i.e., the minimum of one day apart) is 307.

As a result we calculate the time-to-warning as the mean time between *clusters* of warnings. More precisely, we compute the average of the times between warnings that exceed a quarter year. We show the result for $n \in \{2,3\}$ and $s \in \{250, 300, 350, ..., 1000\}$ in Figure S14. For each value of n there is a roughly linear relation between the sliding window size s and the time-to-warning. The times between warnings vary roughly between one and ten years, which is a magnitude that is consistent with our intuition that financial crises happen in the order of years, but not tens of years.

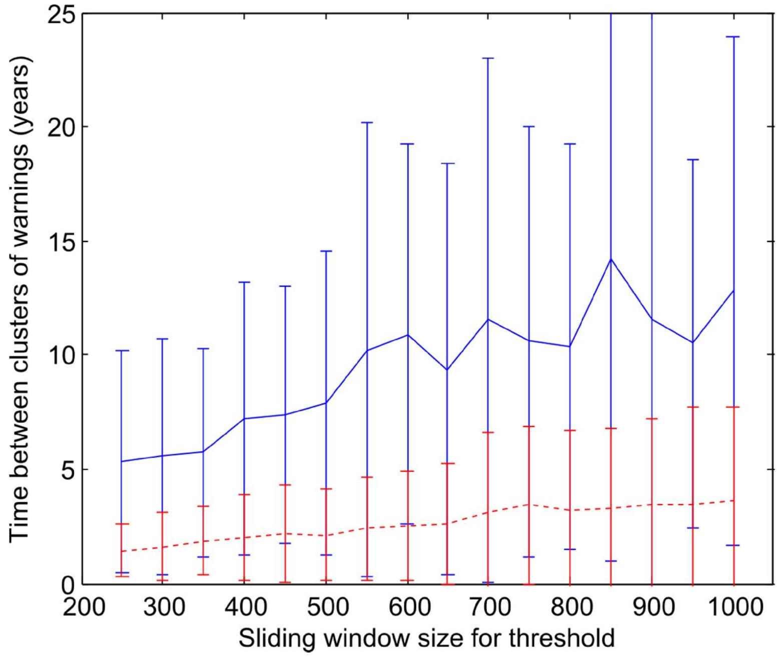

Figure S15: The expected number of years between two clusters of warnings induced by the simulated AR(2) IDL indicator. The red dashed line corresponds to $n = 2$ and the blue sold line to $n = 3$. Each error bar is one standard deviation.

Selecting the appropriate mean time-to-warning is a matter of the magnitude of market crashes that one is interested in. Here we use statistics gathered by Laeven and Valencia[24] from the Systemic Banking Crises Database to find a suitable value. This database records market crashes during 1970--2011 in three forms of increasing severity: single crises, twin crises, and triple crises. A triple crisis is a banking crisis that coincides with a currency crises as well as a sovereign crisis; a twin crisis is any pair of these crises. We choose the middle way between twins and triplets as our magnitude of interest, so we compute the average of their frequencies and select it as our time-to-warning. According to Laeven and Valencia[24], the global frequency of twin crises is once in 0.54 years and the frequency of triple crises is once in 5.13 years. The average between the two is a time between crises of 2.83 years.

This value falls in the range for $n = 2$ in Figure S15. To find the corresponding value for s we perform a linear regression on the mean values of the time-to-warning curve, in order to filter out the fluctuations. For the $\mathrm{AR}(2)$-generated IDL indicator we find an optimal value $s = 691.44$; if instead we use $\mathrm{AR}(3)$ to generate the IDL then the result is $s = 709.46$.

We repeat the procedure for the critical slowing down indicator shown in Figure 2 in the main text. Here we use AR(1) and find $s = 701.71$ (EUR) and $s = 695.40$ (USD) for $w = 1000$.

Since all these values are similar we round them up to the nearest quarter year which is 750 trade days, corresponding to three calendar years.

In conclusion, we find the parameter values $n = 2$ and $s = 750$ to generate clusters of warnings about once in three years, which is a tradeoff between the frequency of the most severe crises (triplets) and the medium crises (twins).

S8 The data

We make all data sets publicly available online in the SI.

S9 The program code

We have performed all analyses and created all plots using Matlab 2010b. All necessary source code to reproduce our results is available online in the SI.

S10 References

1. Cover, T. M. & Thomas, J. A. *Elements of information theory*. **6,** (Wiley-Interscience, 1991).
2. James, R. G., Ellison, C. J. & Crutchfield, J. P. Anatomy of a Bit: Information in a Time Series Observation. *Chaos* **21,** 15 (2011).
3. Dakos, V., van Nes, E., Donangelo, R., Fort, H. & Scheffer, M. Spatial correlation as leading indicator of catastrophic shifts. *Theoretical Ecology* **3,** 163–174 (2010).
4. Lenton, T. M., Livina, V. N., Dakos, V., van Nes, E. H. & Scheffer, M. Early warning of climate tipping points from critical slowing down: comparing methods to improve robustness. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **370,** 1185 –1204 (2012).
5. Donangelo, R., Fort, H., Dakos, V., Scheffer, M. & Nes, E. H. Early warnings for catastrophic shifts in ecosystems: Comparison between spatial and temporal indicators. (2010).
6. Bailey, R. M. Spatial and temporal signatures of fragility and threshold proximity in modelled semi-arid vegetation. *Proceedings of the Royal Society B: Biological Sciences* **278,**

1064 –1071 (2011).
7. Guttal, V. & Jayaprakash, C. Spatial variance and spatial skewness: leading indicators of regime shifts in spatial ecological systems. *Theoretical Ecology* **2,** 3–12 (2009).
8. Reshef, D. N. *et al.* Detecting Novel Associations in Large Data Sets. *Science* **334,** 1518–1524 (2011).
9. Cellucci, C. J., Albano, A. M. & Rapp, P. E. Statistical validation of mutual information calculations: Comparison of alternative numerical algorithms. *Phys. Rev. E* **71,** 066208 (2005).
10. Rissanen, J., Speed, T. P. & Yu, B. Density estimation by stochastic complexity. *Information Theory, IEEE Transactions on* **38,** 315–323 (1992).
11. Mosteller, F. & Tukey, J. W. *Data Analysis and Regression*. (Pearson, 1977).
12. Bendat, J. S. & Piersol, A. G. *Random Data: Analysis and Measurement Procedures*. (Wiley, 2010).
13. Terrell, G. R. & Scott, D. W. Oversmoothed Nonparametric Density Estimates. *Journal of the American Statistical Association* **80,** 209–214 (1985).
14. Wissel, C. A universal law of the characteristic return time near thresholds. *Oecologia* **65,** 101–107 (1984).
15. Dakos, V. *et al.* Slowing down as an early warning signal for abrupt climate change. *Proc. Natl. Acad. Sci. U.S.A.* **105,** 14308–14312 (2008).
16. Dakos, V., Kéfi, S., Rietkerk, M., Nes, E. H. van & Scheffer, M. Slowing Down in Spatially Patterned Ecosystems at the Brink of Collapse. *The American Naturalist* **177,** E153–E166 (2011).
17. Scheffer, M. *et al.* Early-warning signals for critical transitions. *Nature* **461,** 53–59 (2009).
18. Drake, J. M. & Griffen, B. D. Early warning signals of extinction in deteriorating environments. *Nature* **467,** 456–459 (2010).
19. Rajdeep Sengupta & Yu Man Tam. The LIBOR-OIS Spread as a Summary Indicator. *Economic Synopses* (2008).
20. Estrella, A. & Mishkin, F. S. Predicting U.S. Recessions: Financial Variables as Leading Indicators. *Rev. Econ. Stat.* **80,** 45–61 (1998).
21. Riccardo Rebonato. *Volatility and Correlation: The Perfect Hedger and the Fox*. (Wiley, 2004).
22. Gabaix, X., Gopikrishnan, P., Plerou, V. & Stanley, H. E. A theory of power-law distributions in financial market fluctuations. *Nature* **423,** 267–270 (2003).
23. Preis, T., Schneider, J. J. & Stanley, H. E. Switching processes in financial markets. *Proceedings of the National Academy of Sciences* **108,** 7674–7678 (2011).
24. Valencia, F. & Laeven, L. *Systemic Banking Crises Database: An Update*. (International Monetary Fund, 2012).