Article

Noise Reduction in a Reputation Index

Peter Mitic [1,2,3]

[1]   Santander UK, 2 Triton Square, Regent's Place, London NW1 3AN, UK; peter.mitic@santandergcb.com;
Tel.: +44-207-576-1234
[2]   Department of Computer Science, University College London, Gower Street, London WC1E 6BT, UK
[3]   Laboratoire d'Excellence sur la Régulation Financière (LabEx ReFi), 75272 Paris, France

Abstract: Assuming that a time series incorporates "signal" and "noise" components, we propose a method to estimate the extent of the "noise" component by considering the smoothing properties of the state-space of the time series. A mild degree of smoothing in the state-space, applied using a Kalman filter, allows for noise estimation arising from the measurement process. It is particularly suited in the context of a reputation index, because small amounts of noise can easily mask more significant effects. Adjusting the state-space noise measurement parameter leads to a limiting smoothing situation, from which the extent of noise can be estimated. The results indicate that noise constitutes approximately 10% of the raw signal: approximately 40 decibels. A comparison with low pass filter methods (Butterworth in particular) is made, although low pass filters are more suitable for assessing total signal noise.

1. Introduction

Practical ways to measure reputation have advanced rapidly in the past few years with the increased use of internet data feeds. As with any measurement that involves a physical apparatus, error is induced in such a measurement by the measurement process itself. The purpose of this paper is to estimate the extent of the error induced by the process of collecting and analysing sentiment data to produce a reputational time series. It is often argued that it is important for an organisation to maintain a "good reputation", but two questions are commonly left unanswered. The first is the question of what precisely is meant by the term "reputation". The second is how to measure it. For an example of this approach, see Deloitte (2016). A slight advance is exemplified by Cole, in which reputation is determined by survey, and reputational effects are noted using changes in balance sheet items. The first of these questions has been addressed in detail in Mitic (2017a). The second was answered in Mitic (2017b). To provide a context for the main discussion of this paper, the results of those two papers will be summarised below. Given that reputation can be measured, the measure becomes meaningful in commercial terms only when expressed in monetary terms. A conversion in terms of annual company sales and profit was estimated in Mitic (2017b), thereby justifying the effort and expenditure that is often directed at maintaining a positive reputation. The error induced in a measure of reputation by the process of measurement therefore translates into an error in associated monetary risk. The amount of this error is estimated in the final part of this paper. Three methods for estimating the amount of noise in reputational time series are described and compared. The State-Space (Kalman filter) method is the principal method, and it is compared with a simple low pass filter (*Moving Average*), and a more complex low pass filter (*Butterworth*). Such low pass filters are more suited to assessing the total signal noise because they do not have a dedicated parameter that controls measurement noise, as does the Kalman filter. Nevertheless, low pass filters give an overall indication of signal noise. The

terms *time series* and *signal* are used interchageably in this paper, the terms being used to mean the same thing in the worlds of statistics and signal processing, respectively.

2. Reputation and Its Measurement

This section contains a summary of the main points in Mitic (2017a) and Mitic (2017b). The process of reputation measurement is a relatively recent development, made possible by exploiting internet connectivity and advances in sentiment analysis.

2.1. Definitions for Sentiment and Reputation

Consider a set of agents (individuals or groups) that comment on a target organisation G at time t. Those comments can be analysed and each can be given a score in the range $[-1, 1]$. See Liu (2015) for a full discussion of how such a score may be derived. A score near 1 means that the agent who expressed the comment is very favourably disposed towards G. A score near -1 means exactly the opposite, and a score of 0 means that the agent's position with respect to G is neutral. The range $[-1, 1]$ is arbitrary, but it makes intuitive sense to cast the score in terms of positive and negative numbers.

To be more precise, let H (for Holder) denote an agent who makes a comment at time t, and denote the corresponding score by $s_{i;H}$, where i is a unique identifier for the comment. Next, consider a collection of scored comments $\{s_{i;H}\}_{t \in T}^{i \in I}$, in which T is an indexing set for times and I is an indexing set for unique identifiers. The reputation of organisation G at time t, $R_G(t)$, can then be defined as a weighted average of the scores in such a collection. To this end, let w_i be the weight of comment i. Then, $R_G(t)$, can be expressed as:

$$R_G(t) = \frac{\sum_{i=1}^{n} w_i s_{i;H}}{\sum_{i=1}^{n} w_i}. \tag{1}$$

The weights w_i are arbitrary, but represent some appropriate characteristic of the holder of the corresponding score $s_{i;H}$. For example, w_i could represent the influence of H. A high (positive) weight indicates a very influential holder (for example, a national newspaper) and low (positive) weight indicates an uninfluential holder (for example, a Twitter user with only a few followers). For example, a Twitter user with 100 or fewer followers is assigned weight 1. A Twitter user with 1 million or more followers is assigned weight 10. Other uses are weighted on a linear scale based on the two fixed points defined by those users. The number of followers for each user under consideration is determined dynamically. Similarly, for traditional media, the weights are based on circulation and viewer figures for newspapers and broadcasts, respectively. Note that Equation (1) applies only at a single time t. It is only a single snapshot. More generally, the reputation of G is formally defined as a time series, $\hat{R}_G$ (Equation (2)).

$$\hat{R}_G = \{R_G(t)\}_{t \in T}. \tag{2}$$

We stress that Equation (2) does not account for opinion formulation on the basis of only a few (perhaps only one) received comments. This is a form of "reputation" that sometimes appears in informal discourse, and is not appropriate for the analysis in this paper.

2.2. Sentiment Procurement

Figure 1 shows an overview of the data mining and analytical stages in sentiment procurement. Proceeding from left to tight in the figure, the stages are:

1. Receive "comments" electronically from opinion holders, H, that convey sentiment with respect to a target G from public sources (news channels, social media, etc.).
2. Sentiment analysis for each comment, to give a sentiment "score", nominally in the range $[-1, 1]$.
3. For each comment, define a weight (e.g., to reflect the influence of the opinion holder of the comment).
4. Compose a reputation index applicable at a particular time using all the sentiment scores received in a given time period (such as one day). This is Equation (1).

5. Accumulate successive time-based reputation indexes to form the reputation time series, $\hat{R}_G$ of Equation (2).

The result is a single number per time period (in practice daily) that measures the reputation of the target organisation, G. The approach indicated above was pioneered by the London business intelligence consultancy Alva Group (www.alva-group.com). The precise calculation of the elements $R_G(t)$ in Equation (1) is more complex than indicated here, and is not publicly available.
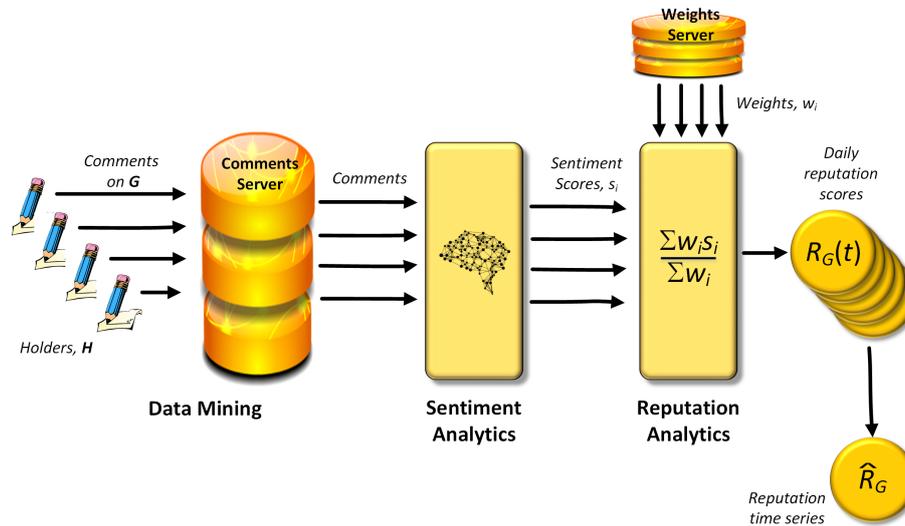

Figure 1. The stages in the process of reputation measurement.

The short series, S, i.e., Equation (3), shows a few examples of the result of the process shown in Figure 1 (Ford Motors, starting on 1 July 2015, adapted from data from Alva Group). Positive and negative numbers correspond ,respectively, to positive and negative sentiment.

$$S = [-0.1815, -0.0036, -0.0967, 0.2937, 0.0135, -0.0316] \tag{3}$$

The entries in Equation (3) show some numbers with relatively high absolute values, and a few that are close to zero. The latter are the subject of this paper. This point will be explored in more detail shortly in this paper.

2.2.1. Reputation and Financial Risk

Basic definitions of sentiment and reputation were given in Mitic (2017a). The essential premise of that paper is that reputation can be expressed in terms of monetary units (GBP, USD, EUR, etc.). The finding differentiated between three operational cases. The first is a "business as usual" (*BAU*) context, in which there are no major reputational shocks. The second is the "stressed" context, in which a significant reputational event results in a noticeable change in the values in the reputation index following the event. The third is the "super-stressed" context, in which a very significant reputational event results in a prolonged and deep change to the profile of the reputation index. An example of a "super-stressed" event is the Volkswagen emissions revelation in September 2015. See Jung and Park (2016) or the wiki entry (https://en.wikipedia.org/wiki/Volkswagen_emissions_scandal) for an account of the history of this affair. The conversions from accumulated annual reputation to profit after tax are shown in Table 1.

Table 1. Effect of annual reputation change on profit after tax.

Context	Positive Reputation	Negative Reputation
BAU	0.6%	−1.1%
Stressed	2.3%	−2.7%
Super-Stressed	5%	−15%

In the analysis that follows, a convenient risk metric, and also a telling visual illustration, is to calculate the cumulative reputation, $\tilde{R}_G$, which is the set of cumulative sums of the reputation scores $R_G(t)$ in $\tilde{R}_G$. With the indexing set T for times t:

$$\tilde{R}_G = \{\sum_{s=1}^{t} R_G(s)\}_{t \in T}. \tag{4}$$

The way in which $\tilde{R}_G$ changes with small changes in particular values of $R_G(t)$ can be expressed in monetary terms. This is the methodology proposed in Section 3 to assess the effect of measurement error in values of $R_G(t)$.

2.2.2. Reputation Measurement Error

There is an induced error in measuring reputation in the way described above due to data collection and sentiment analysis of that data. A small error or change in the detailed sentiment assessment for a component of a term $R_G(t)$ can reverse the sign of $R_G(t)$, provided that the value of $R_G(t)$ is very marginally positive or very marginally negative. The purpose of this study is to assess the extent of this error and to calculate how much it is worth in monetary terms. This permits an error bound on $R_G(t)$ to be calculated in the following way. First assume that a set of very marginal negative reputation values should, in fact, have been positive. Changing their signs from negative to positive then increases the proportion of positive sentiment scores that enter the calculation of $R_G(t)$. Similarly, a set of sentiment scores, incorrectly recorded as positive, could have their signs changed to negative, and thereafter be treated as negative scores. This would then increase the proportion of negative components values in $R_G(t)$. The process of transferring negative sentiment values to positive, and positive to negative, followed by a recalculation of $R_G(t)$ in the two cases, allows us to estimate an upper and a lower bound on the value of $R_G(t)$ induced by the measurement process. The monetary value of those error bounds can then be estimated.

As a simple example, consider the second reputation score, −0.0036, in S (Equation (3)), which is close to zero. Suppose that the weighted components of that sum are 0.245, 0.109, −0.078, −0.27, −0.0122. Suppose, further, that a change to the sentiment analysis algorithm results in a small change to one of these components. For example, the first component, 0.245 changes to 0.260. That is enough to change the sum of the five components from negative to a positive value, 0.0088, thereby reversing the sign and sense of the reputation score. Alternatively, suppose that a change to the data collection procedure results in an additional weighted component, 0.007. When the six components are summed, the result is 0.0008, again reversing the overall sign and sense of the reputation score.

Following this numerical illustration, we now discuss the detailed algebraic implementation of a method to calculate measurement error bounds.

3. Methodology

The methodology used to assess the inherent measurement error described in the previous section is to apply a smoothing filter to the times series $\hat{R}_G$ to isolate, objectively, a set of positive and negative values of $R_G(t)$ that are in the neighbourhood of zero. The upper bound is then obtained by mapping the negative values in that neighbourhood to a positive value, thereby boosting the cumulative reputation, $\tilde{R}_G$. Similarly, the lower bound is obtained by mapping the positive values

in that neighbourhood to a negative value, thereby reducing the cumulative reputation. State space smoothing is the preferred method used to achieve the above objectives. The reasons are given in the next sub-section. The following subsections contain a summary of the state-space methodology, and the specialisations of it applicable to the task in hand. The State-Space method is one of many methods that filter parts of a signal. We associate noise with the high frequency parts of a signal, and aim to remove that noise using the filter. The *Kalman* filter is associated with the State-Space method. More generally, we also consider removal of high frequency noise using other *low pass*. They are designed to transmit the low frequency parts of the signal, and to block the high frequency parts. In general we are less confident in using low pass filtering for our purpose because they have no specific parameter to assess measurement noise.

3.1. Nomenclature

For the purposes of this paper, a time series received from a source without further processing or analysis will be referred to as an *Observed Signal*. Such an *Observed Signal* comprises two components: an *Unobserved Signal*, which is the part we would like to discover, and *Noise*. The *Unobserved Signal* and the *Observed Signal* at time t will be denoted by the symbols x_t and y_t, respectively. In the uni-dimensional state-space analysis that follows, both y_t and x_t are scalars, but more generally they are vectors.

3.2. The State-Space and the Kalman Filter

State space analysis, otherwise known as "dynamic linear modelling", originates from the work of Kalman (1960) in the context of tracking the position (i.e., the *state*) of a spacecraft with location x_t, given a noisy location measurement y_t. x_t is termed the *state vector*. It contains all available information to describe the investigated system, and is usually multidimensional. The measured vector, y_t, represents observations related to the state vector. y_t is often of lower dimension than the state vector. The idea that some quantity has a value which we can estimate by measuring is central to the argument in this paper. Measuring reputation by collecting expressed sentiments electronically is not exact. A certain amount of expressed sentiment constitutes "noise" in the sense that the sentiment is expressed routinely such that it does not reflect any significant event. Chapter 6 in Shumway and Stoffer (2017) is a good account of the Kalman filter in the context of state-space analysis, and the summary that follows is based on it. The Kalman filter can be used for any of three purposes: prediction, filtering and smoothing. We exploit the latter here. The name *filter* should be thought of as a filtering out of noise (i.e., noise *reduction*, not elimination). A useful phrase summarising Kalman filtering is:

> ***Noisy signal in: less noisy signal out***

The Kalman filter technique is used in the context of reputational analysis for three main reasons:

- Kalman filtering incorporates a specific smoothing parameter which can be used to assess measurement noise. Other smoothing methodologies (e.g., Loess filter and moving average) do not have such a specific relationship with measurement.
- Construction of the reputation index (specifically resetting to an initial value every period) produces a signal for which a static model with noise is appropriate: there is necessarily no trend and no seasonality. A particularly simple version of the Kalman model is available for precisely this situation.
- Very little lag is introduced by the smoothing process.
- The Kalman filter (state-space) method makes direct use of the distribution of reputation scores, which can be modelled successfully by a Normal distribution.
- The Kalman filter methodology assumes that the time series being analysed is stationary. All of the reputational time series in this study are stationary, as demonstrated by examining their

frequency spectra. It is expected that others will also be stationary unless they are driven by seasonal events.

3.3. Kalman Filter Fundmentals

Let the current unobserved state of a quantity x at time t be x_t, where t varies discretely in steps of 1 between 0 and a maximum value τ. In one dimension, this is the single state variable. Then, in terms of the state at the prior time slot $t-1$,

$$x_t = F_t x_{t-1} + B_t u_t + w_t, \tag{5}$$

in which u_t is a control variable that can influence x_t and B_t is, in general, a control matrix which maps controls to state variables. w_t is a stochastic error which is usually taken as $N(0, Q)$, where Q is a correlation matrix measurement of volatility of the state variable itself. Ultimately, we aim to calculate an estimator, $\hat{x}_t$, of x_t. The observed variable y_t is related to x_t by Equation (6), in which u_t also appears with another control matrix C_t. The matrix H_t is known as the *observation matrix* and represents a measurement scaling. There is another stochastic error term v_t which is usually taken as $N(0, R)$, where R is a measurement error matrix.

$$y_t = H_t x_t + C_t u_t + v_t \tag{6}$$

In the uni-dimensional case, all matrices in Equation (6) are 1-by-1, so they are scalars.

The Kalman analysis continues by defining a *predict* and an *update* stage. Let $\hat{x}_{t|t}$ denote the estimate of state x_t at time t. In addition, let $P_{t|t}$ be the variance matrix of the error $x_t - \hat{x}_{t|t}$. The goal is to minimize $P_{t|t}$. The prediction stages for $\hat{x}$ and P are (the superscript T denotes transpose):

$$\hat{x}_{t|t-1} = F_t \hat{x}_{t-1|t-1} + B_t u_t; \quad 1 \leq t \leq \tau. \tag{7}$$

$$P_{t|t-1} = F_t P_{t-1|t-1} F_t^T + Q_t; \quad 1 \leq t \leq \tau. \tag{8}$$

The *predict* stage is initialised with known values x_0 for $\hat{x}_{0|0}$, and P_0 for $P_{0|0}$. The corresponding update stage is given in terms of a measurement residual vector $\tilde{z}_t$, a residual variance matrix S_t, and an innovation part which gives the transition of $\hat{x}_{t|t-1}$ to $\hat{x}_{t|t}$ in terms of a *Kalman gain* matrix, K_t:

$$\begin{aligned} \tilde{z}_t &= y_t - H_t \hat{x}_{t|t-1}. \\ S_t &= H_t P_{t|t-1} H_t^T + R_t. \\ \hat{x}_{t|t} &= \hat{x}_{t|t-1} + K_t \tilde{z}_t. \end{aligned} \tag{9}$$

There is an optimal gain Kalman Filter, defined by

$$K_t = P_{t|t-1} H_t^T S_t^{-1}. \tag{10}$$

With this K_t, the update equations resolve to:

$$\begin{aligned} \hat{x}_{t|t} &= \hat{x}_{t|t-1} + K_t (y_t - H_t \hat{x}_{t|t-1}). \\ K_t &= P_{t|t-1} H_t^T (H_t P_{t|t-1} H_t^T + R_t)^{-1}. \\ P_{t|t} &= (I - K_t H_t) P_{t|t-1}. \end{aligned} \tag{11}$$

The three *predict* Equations (9) and the three *update* Equations (11) define the basis of the special case that follows.

3.3.1. Kalman Smoothing: The Static Model Specialisation

In this section, we modify the *predict* and *update* equations such that they apply to the case where there is effectively a constant noisy signal. In Mitic (2017a) the idea of a standardised reputation score was introduced and developed (Subsection 3.3.1 in this paper has a summary). A daily score in the range $[-1, 1]$ is produced, the range $(0, 1]$ corresponding to positive sentiment and the range $[-1, 0)$ corresponding to negative sentiment. A score of 0 represents completely neutral sentiment, and a score in a neighbourhood of 0 is the noise that we wish to estimate. This score is the observed signal y_t, time t referring to a daily measurement. Details of the physical method of sentiment procurement were also given in Subsection 3.3.1, and a fuller account may be found in Mitic (2017b). In that paper, and also in Subsection 2.1 of this paper, the formal time series definition of the reputation of an organisation G at time t is written as $R_G(t)$, but the notation y_t serves for the current purpose because it is consistent with much of the notation on state-space analysis in the literature. The methodology is such that the neutral score is the starting point each day.

As such, we regard the resulting reputational time series as measuring deviations from a constant value (namely the neutral score), with, necessarily, no trend and no seasonality. This condition allows us to set the 1-by-1 transition matrix in Equation (7) to 1: $F_t = 1$; $0 \leq t \leq \tau$. In addition we assume there is no control term, so $u_t = 0$. The matrix H represents a scaling with respect to the measurement, and is also set to 1 since we make no measurement scaling. The process is a scalar, so P is a scalar. Process noise Q is assumed to be a constant scalar with value q. The same applies for measurement noise, so we set R to a constant scalar value r. The *predict* Equations (9) and the *update* Equations (11) then reduce to the following equations for $1 \leq t \leq \tau$. The smoothing case is for $0 \leq t < \tau$. The case $t = \tau$ corresponds to filtering and if $t > \tau$, the Kalman process is in predictive mode.

Predict:

$$
\begin{aligned}
\hat{x}_{t|t-1} &= \hat{x}_{t-1|t-1}; \quad 1 \leq t \leq \tau. \\
P_{t|t-1} &= P_{t-1|t-1} + q; \quad 1 \leq t \leq \tau. \\
\hat{x}_{0|0} &= x_0. \\
P_{0|0} &= P_0
\end{aligned}
\tag{12}
$$

Update:

$$
\begin{aligned}
\hat{x}_{t|t} &= \hat{x}_{t|t-1} + K_t(y_t - \hat{x}_{t|t-1}); \quad 1 \leq t \leq \tau. \\
K_t &= P_{t|t-1}(P_{t|t-1} + r)^{-1}; \quad 1 \leq t \leq \tau. \\
P_{t|t} &= (1 - K_t)P_{t|t-1}; \quad 1 \leq t \leq \tau.
\end{aligned}
\tag{13}
$$

3.3.2. Normality Assumption of State-Space Analysis

The state-space methodology, summarised in terms of the Kalman filter, by Equations (9) and (11), and the specialisation in Equations (12) and (13) depend on a normality assumption. Specifically, the elements of the input time series, y_t, must be normally distributed. State space is a linear representation of the dynamic behaviour of y_t, so if y_t is normally distributed, so is its state-space formulation, x_t (Hamilton 1986). Further, the Kalman filter algorithm assumes that the state transitions are linear, and that the error terms when fitting to data are normally distributed. See, for example, Cosma and Evers (2010) or Shumway and Stoffer (2017) for a full discussion of this point. Calculation of confidence limits in the analysis that follows depends on this assumption of normality.

If the input time series y_t is not normally distributed, it is generally possible to apply a transformation such that the transformed series is normally distributed. Generally, power transformations are successful, and in particular the Box-Cox transformation (Equation (14), which depends on selecting a suitable parameter λ. Using Box-Cox, each element y of y_t is mapped to y', where

$$
y' = (y^\lambda - 1)/\lambda.
\tag{14}
$$

Alternatively, *log* transformations are often successful in transforming to normal distributions. Many of the input time series used in this analysis have outliers that indicate non-normality. Therefore a Box-Cox transformation has been applied to normalise them sufficiently to permit calculation of confidence limits.

3.3.3. Kalman Filter Noise Estimation

Equations (12) and (13) can now be applied to estimate the noise component in a signal. The first stage is to fix the volatility parameters q and r. Parameter q is taken as the standard deviation of the observed signal. It is assumed to be an intrinsic property of the observed signal, and is therefore fixed. Parameter r represents measurement noise, and is determined by considering a range of appropriate values and choosing one that results in maximum noise, as determined by the method outlined below. Once a value for r has been found, the noise in the observed signal can be isolated using the confidence interval, J, given by Equation (15). In (15), m_x and s_x are the mean and standard deviation respectively of the unobserved (i.e., smoothed) signal x_t, and $z_{0.99}$ is the 99% 2-tail critical value on the standard Normal distribution. In addition, recall that τ is the number of elements in both observed and unobserved time series. The confidence interval is based on the assumed Normal distribution for the unobserved signal once it has been modified by Box-Cox.

$$J = \left(m_x - z_{0.99}\frac{s_x}{\sqrt{\tau}}, \quad m_x + z_{0.99}\frac{s_x}{\sqrt{\tau}} \right) \tag{15}$$

The values of m_x and s_x vary with the value of the parameter r. For each r, the number of unobserved signals within the confidence interval J can be identified and counted. Writing $\{y_t \in J | r\}$ and $\{y_t \notin J | r\}$ for the sets of observed signals in J and the observed signals not in J, respectively (so that $\tau = count\{y_t \in J | r\} + count\{y_t \notin J | r\}$), the percentage of noise in the unobserved signal, P, is then given by Equation (16). This percentage is taken to be the percentage of noise in the observed signal.

$$P = 100\,\frac{count\{y_t \in J | r\}}{\tau} \tag{16}$$

3.4. Low Pass Filter Noise Estimation

The action of a low pass filter is to transmit signals with a frequency lower than a cutoff frequency, and to reject signals with frequencies higher than the cutoff. In signal processing, the word *attenuate* is used is used in place of "reject". Using a low pass filter on the reputational signals considered in the paper is based on the assumption that the informative parts of the signal have the more extreme values (near +1 or −1). They correspond to the low frequency part of the signal, simply because there are more of them. Noise is confined to values near to zero. There are many more of them, and we associate those values with high frequency. In principle, it is difficult to link the concept of *frequency* in the context of a reputation signal with the way "frequency" is used in the context of an audio signal. Audio signals are generally treated as superpositions of sinusoidal sub-signals in the form of a Fourier series. This approach has not been hitherto been applied to a reputation signal. We attempt to estimate the amount of noise by filtering out the high frequency noise. In general, low pass filters have only one parameter available for controlling which frequencies are filtered out. Therefore, they filter out all noise, not just measurement noise. They are therefore not as suitable as a means to assess measurement noise, but it is interesting to see the results of using this type of filter. Smith (2007) (Chapter 1) has a simple account of how a low pass filter is intended to work, and Hamming (1989) has a more rigourous account. Practical guidance on programming digital filters using *Matlab* may be found in Jackson (1996). In principle, any method of smoothing a time series can be used as a filter. None has a dedicated parameter for measurement noise, so they are not as suitable as the Kalman filter for our purposes. The following sections concentrate on two: *Butterworth*, (Butterworth (1930)), because the low pass property is easy to define, and *Moving Average*, as a simple comparison.

The overall way in which a digital filter works is to define a new series by applying a transformation T_t to y_t (Equation (17)). Thus, if the filtered signal is x_t,

$$x_t = \{T(y_t)\}_{t \in [1..n]}. \tag{17}$$

Figure 2 illustrates the use of a low pass filter. The solid line shows an ideal state where frequencies less than a cutoff frequency f_c are transmitted, and where no frequencies greater than or equal to f_c are transmitted. The vertical axis shows the *gain*, which is the ratio of the amplitude of the output signal to the input signal. The "ideal state" filter is shown in red. The horizontal portion corresponding to frequencies f where $0 < f < f_c$, and is also shown. All frequencies in that range are transmitted with no degradation in amplitude, shown by gain = 1. The gain for $f \geq f_c$ is zero because, ideally, no signal is transmitted in that range. The "non-ideal state" filter is shown in blue. The gain is sub-optimal due to signal degradation, the cutoff at frequency f_c is not sharp, and high frequencies are not completely removed. In the context of a reputation signal, frequency is ill-defined. It is replaced by an appropriate parameter of the method used to assess the noise level. The green profile shows a frequency response where there is no clear cutoff. This shows a response that often occurs in practice: see Subsection 3.4.2.
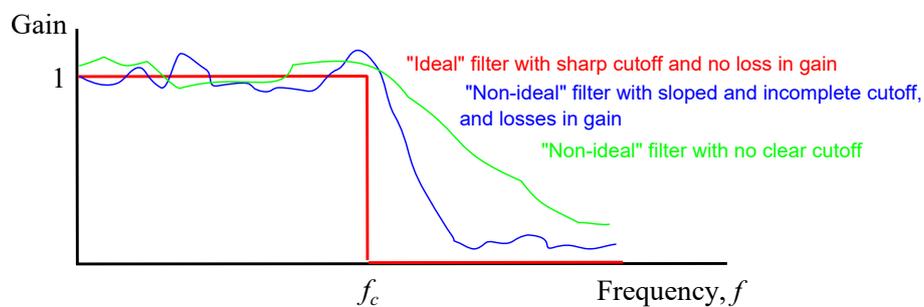

Figure 2. Ideal (red) and non-ideal (blue) filters, showing gains as a function of Frequency.

3.4.1. The Power Spectral Density (PSD)

The PSD (sometimes called the Power Spectrum) of a signal shows the power in the signal as a function of frequency, per unit frequency. For a discrete signal, power is measured as the square of the signal amplitude divided by the number of signals. A detailed theoretical basis may be found in Stoica (2005). For a simpler treatment, see Cerna (2009). The PSD calculation starts by calculating the Discrete FourierTransform (DFT) of the signal, y_t (Equation (18), in which f is a frequency and n is the number of elements in the series $\{y_t\}$).

$$F(f) = \sum_{t=0}^{n-1} y_t e^{-2\pi i f t/n} \tag{18}$$

The DFT is a complex number for every f, and its amplitude is given by $|F(f)|$. Therefore, the PSD is given by the series in Equation (19), in which ϕ is a set of frequencies.

$$PSD = \{(F(f))^2/n\}, f \in \phi \tag{19}$$

The PSD is used in our context to examine the effectiveness of the filtering methodology used. The shape of the PSD is sometimes indicative of a clear noise cutoff level, such that we can regard frequencies above that cutoff as "noise".

3.4.2. Low Pass Filter Procedures

The steps in the general approach used in applying a low pass filter to a reputation signal are as follows. This approach makes use of the PSD to find a noise cutoff frequency.

1. Obtain the PSD for the input signal under consideration
2. Determine a cutoff frequency, f_c, by finding the point at which the power spectrum stabilises. All power spectra considered have the property that the power density stabilises at some fraction of the total number of ordinates in the spectrum. Most reputation power spectra exhibit shape frequency profiles similar to the profile illustrated in Subsection 4.2.2. The plot shows a limiting value at frequency 0.7. There is one exception where the spectrum indicated periodicity with period of 100 days (organisation NW in Table 4). This periodicity was likely due to chance, as no reason for it was apparent and it was not observed for other organisation. There is a discussion of special treatment of two other organisations in Subsection 4.3.
3. Apply a filter to the input signal y_t using the cutoff frequency f_c to obtain the output signal x_t
4. Obtain the PSD for the output signal, x_t. If there are n PSD components $p_1, p_2, ..., p_n$, denote the normalised cumulative sums of the square of those components by $S_t = \sum_{i=1}^{t} p_i^2$ (squaring emphasises any distinction between low and high frequencies).
5. Estimate the noise level, P (expressed as a percentage) from Equation (20). In Equation (20), $\#[S_t > q]$ denotes the number of S_t entries that are greater than q.

$$P = 100 \frac{\#[S_t > q]}{n} \tag{20}$$

Two types of filter will be considered for this analysis. The Butterworth filter aims to achieve an "ideal" filter profile (as in Figure 2). The Moving Average filter is discussed as a comparison.

3.4.3. Moving Average Filter

The Moving Average filter is a very simple form of low pass filter, and is generally used for smoothing time series. When used for filtering, the only parameter that can be used is the number of terms in the moving average. Moving averages introduce lag, so if a very fast response is needed, this may not be an ideal filter method. However, it is very simple to define and implement. If y_t denotes the input time series, the output time series (i.e., the filtered series) of order n, x_t, is given by Equation (21).

$$x_t = \frac{1}{n} \sum_{i=0}^{n-1} y_{t+i} \tag{21}$$

The main use of a Moving Average filter in practice is to reduce random white noise while keeping a sharp "ideal" step response, as in Figure 2. There is a discussion of this filter and variations of it in Smith (1999), Chapter 15.

3.4.4. Butterworth Filter

The Butterworth filter is more complex than the Moving Average filter, and is designed to have as flat a frequency response as possible in the *pass band* (frequencies that should be transmitted), and to block as many frequencies as possible in the *stop band* (frequencies that should not be transmitted). The *pass band* can be defined as a frequency range. The useful *pass band* for the purpose of analysing reputational signals is the range $[0, f_c]$, where f_c defines a maximum low frequency to be transmitted. Frequencies greater than f_c, which represent noise, should be blocked. A full discussion may be found in, for example, Oppenheim (1989). For such a filter, the square of the amplitude of the filtered signal (termed the *Frequency Response*), $H_c(f, f_c, n)$ is given by Equation (22), in which the independent variable is the frequency, and n (an integer greater than 1) is the *order* of the filter. The greater the *order*, the steeper the response at the cutoff point. The *Frequency Response* should be as flat as possible for frequencies less than or equal to f_c. Figure 3 shows three examples. The *Frequency Response* parameter is the only parameter available to control the extent of what is nominally noise. As such the Butterworth filter is less suitable for the purpose of assessing the amount of measurement noise (as opposed to the system noise), but it is worth noting the results of using it as a comparison with the Kalman filter. It should be noted that, in practice, not all frequencies less than or equal to f_c are transmitted, and not

all frequencies greater than f_c are blocked. The greatest frequency transmission errors occur near f_c, and for low values filter orders.

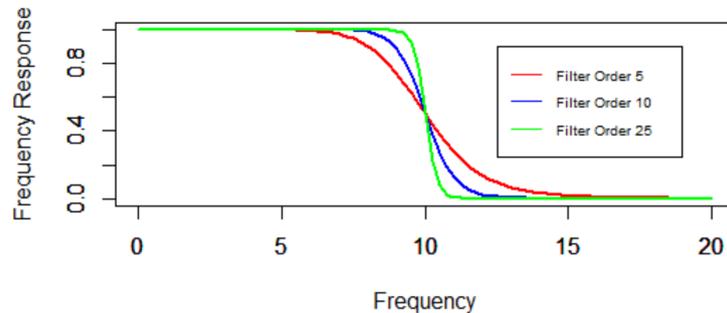$$H_c(f, f_c, n) = \frac{1}{1 + (\frac{f}{f_c})^{2n}} \qquad (22)$$

Figure 3. Examples Butterworth Frequency Responses, critical frequency = 10.

3.5. Signal to Noise Calculation

Expressing the noise level in a time series in terms of the signal to noise (S/N) ratio enables a ready comparison with noise in an audio signal, which may be more familiar to some readers. We therefore include a way to measure the S/N ratio using the terms in Equation (16). There are several established ways to calculate a S/N ratio. See, for example, Smith (1999). A convenient method in this context is to use the square of the ratio of the signal amplitude to the noise amplitude. There is a direct implementation in Matlab: the function *snr*() (see https://uk.mathworks.com/help/signal/ref/snr.html). Thus, if A denotes amplitude,

$$S/N = (\frac{A_{signal}}{A_{noise}})^2 = \frac{\sum_{t=1}^{\tau} \{y_t^2 \notin J | r)\}}{\sum_{t=1}^{\tau} \{y_t^2 \in J | r)\}}. \qquad (23)$$

Expressed in decibels, the S/N ratio is given by $10 log_{10}(S/N)$.

4. Results

In this section, we present numerical illustrations of the theory in Subsections 3.2 to 3.5. Reputation data for ten retail UK banks and six UK motor manufacturers provided the observed (i.e., noisy) time signals y_t, of Equation (6). The data covered the period January 2014 to January 2016. To generate normal distributions for each time series, a Box-Cox transformation was first applied to each. It was sufficient to set a common parameter λ (Equation (14)) to a value derived by aggregating the data from all the time series, and calculating an optimal value. The optimal value was found using package *MASS* in the R statistical language, which exports a function *boxcox*. This function fits a Normal distribution to data that have been transformed using trial values of the Box-Cox parameter λ, using maximum likelihood as the optimisation criterion. The value obtained was $\lambda = 1.846$.

4.1. Kalman Filter Results

The sequence of equations leading to (13) was used to derive values for the unobserved (less noisy) signals, $\hat{x}_{t|t}$. The process noise parameter, q, was set to the standard deviation of the observed signals y_t. Appendix A gives a justification for using this value, and Appendix B gives parameter values and some details of the calculation methodology. The measurement noise parameter, r, was set to values in the range [0.5, 10.5] in steps of 0.1. Values of r less than 0.5 were excluded because they resulted in minimal smoothing, which did not provide an effective assessment of measurement error. With the indicated range it was possible to achieve a moderate degree of smoothing of the observed signal y_t whilst still preserving the profile of the signal. As an example, Figure 4 shows the

first 100 days for both series y_t (the unsmoothed input signal) and $\hat{x}_{t|t}$ (the smoothed output signal) for organisation *FT* when r is set to 0.7, which is the value that results in maximum noise.
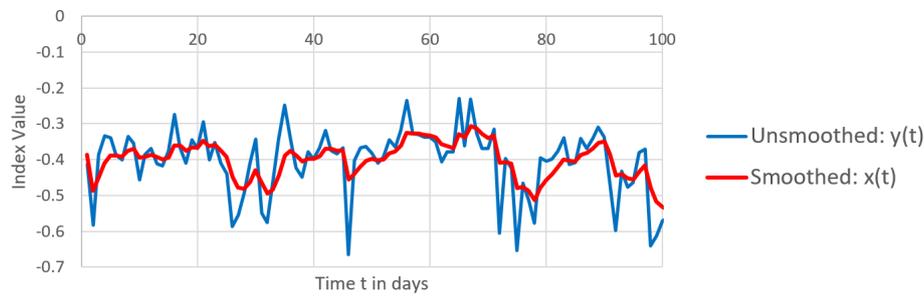

Figure 4. Example Reputation time series (organisation *FT*): measured (observed y_t) and smoothed (unobserved $\hat{x}_{t|t}$), with measurement noise parameter set to 0.7. The plots derive from Box-Cox transformed standard reputation signals.

Figure 4 shows that the degree of smoothing is such that the extremes that exist in the observed signal are damped with a minimal lag. The measurement noise to be estimated is in the region near the horizontal line that represents the mean unobserved signal value. For the complete series the mean is -0.379 (indicating an overall negative sentiment towards *FT*) and for the first 100 entries in the series (shown in Figure 4) the mean is -0.402. Although the theoretical index value can vary between -1 and $+1$, the observed values are clustered closely about the mean. The only way to attain a perfectly positive score of $+1$ or a perfectly negative score of -1 would be for all opinion holders to simultaneously express the same perfectly positive or perfectly negative sentiment. Differing opinions among opinion holders results in a reversion towards a zero score.

Figure 5 shows, for organisation *FT*, how the percentage of noise due to measurement changes as the value of the measurement noise parameter r changes between 0.5 and 10.5. That percentage was calculated from Equation (16) using the confidence interval in Equation (15). Maximum noise is often associated with low values of r. The profile shown is typical.
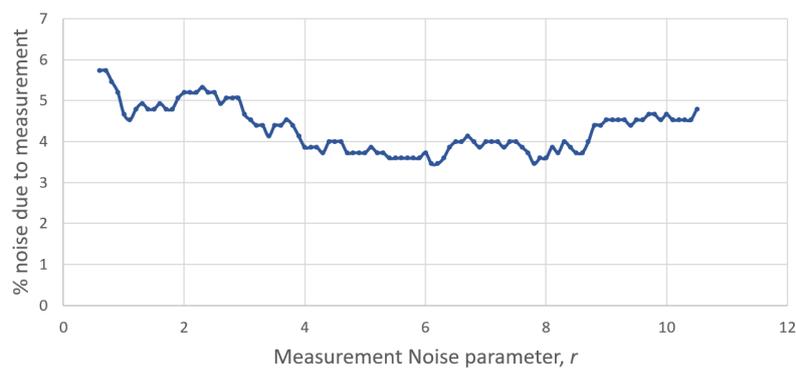

Figure 5. Variation of percentage of signal noise with measurement noise parameter r: organisation *FT*.

Table 2 shows the resulting % noise and S/N ratios in decibels for the organisations considered in this study. The entries in the *Organisation* column are codes to keep the organisations anonymous. The table shows the mean and maximum noise levels as r varies, as well as the S/N values.

An immediate observation from Table 2 is that the noise levels are mostly of the same order, indicating that the results are not merely a feature of the data used. Higher noise levels are largely due to excessive volatility in the observed signal after a major negative reputational hit.

Table 2. Percentage of Noise in Reputation Signals, measured using a Kalman Filter.

Organisation	Mean % Noise	Max. % Noise	S/N in db
BM	7.2	9.5	6.7
FT	4.4	6.3	10.5
FD	6.5	8.7	7.5
HUD	16.1	18.1	9.8
NSS	6.2	7.7	8.5
MB	8.5	10.7	6.2
VOL	4.4	5.2	9.6
BSG	11.1	12.3	9.3
RYB	5.8	8.1	11.7
NW	5.2	7.2	12.1
HX	9.8	13.1	10.7
HS	6.7	7.7	12.5
LL	9.8	11.1	9.2
NT	8.1	12.9	10.6
TB	7.8	10.8	11.3
VG	6.3	9.7	13.3
BB	9.2	10.8	10
DB	4.7	8.3	8.9
Mean	*7.6*	*9.9*	*9.9*

4.2. Low Pass Filter Results

The procedures listed in Subsection 3.4.2, and summarised in Equation (20), were applied separately to each organisation mentioned in the preceding paragraph. Figure 6 is the power spectrum for the *unfiltered* signal of organisation BS in Table 2. It is a typical power spectrum for reputational signals. The plot shows a limiting cutoff value at frequency ∼0.7, indicating a distinction between "signal" and noise at that point. This indicates that, in the *unfiltered* signal, the noise level is about 30%. The following subsections show the effects of applying different low pass filters to this, and other, signals.
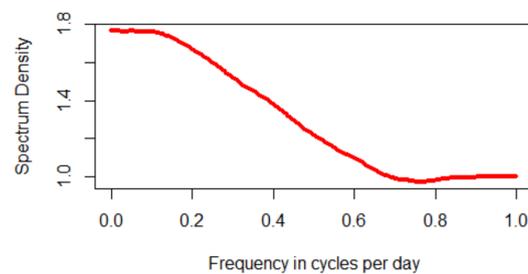

Figure 6. Power Spectrum of the Reputation signal for organisation BS, showing limiting cutoff frequency of approximately 0.7.

4.2.1. Low Pass Filter Results: Moving Average Method

The Moving Average filter is a single parameter method that uses the order of the moving average (i.e., the number of elements over which the average is calculated) to control the separation of signal and noise. The results are shown in Table 3. The corresponding S/N values are calculated using the method of Subsection 3.5.

In general, the mean and maximum noise levels for the Moving Average filter are much higher than the corresponding values for the Kalman filter. This is to be expected because the Kalman filter has been configured to estimate the measurement noise as opposed to the system noise, whereas the Moving Average result incorporates both of these. The S/N ratios are all low compared to reasonable audio signals, indicating that noise is a significant component of the total signal. There is much variation in the results between organisations, so that we can summarise the results by noting the *Mean* entries.

Table 3. Percentage of Noise in Reputation Signals, measured using a Moving Average Filter. The results are means for moving average orders between 2 and 14 (to cover periods of up to two weeks).

Organisation	Mean % Noise	Max % Noise	S/N in db
BW	9.5	25.9	24.1
FT	10.4	24.4	15.3
FD	13.7	41.1	22
HUD	13	29.6	18.5
NSS	18.2	54.3	26.5
MB	13.3	42.6	19.8
VOL	10.1	28.9	7.8
BSG	11.9	29.7	15.8
RYB	6.3	15.4	9.6
NW	3.1	10.2	7.8
HX	10.3	29.7	9.8
HS	8.2	27.1	5.6
LL	11.5	34.1	9.4
NT	19	56	9.7
TB	5.5	14.8	13.5
VG	9.9	28	14.6
BB	8.3	24.2	13.3
DB	14.4	36	21.9
Mean	*10.9*	*30.7*	*14.7*

4.2.2. Low Pass Filter Results: Butterworth Method

The Butterworth filter method uses a single parameter to control the separation of signal and noise: the filter *order*. The mean and maximum noise levels for filter orders between 10 and 25 inclusive was then calculated. These limits provide a reasonably steep cutoff profile (aiming to be similar to the "ideal" filter in Figure 2) without undue complexity. Table 4 shows the results. This table also shows the corresponding S/N values, calculated using the method of Subsection 3.5.

Table 4. Percentage of Noise in Reputation Signals, measured using a Butterworth Filter. The results are means for filter orders between 10 and 25 inclusive.

Organisation	Mean % Noise	Max % Noise	S/N in db
BW	30.7	30.8	13.9
FT	32	32.3	21.9
FD	12.4	12.6	35.4
HUD	24.8	25.5	17.9
NSS	6.2	22.7	66.8
MB	12.1	12.6	31.8
VOL	9.9	13.3	6.9
BSG	25.2	25.8	13.9
RYB	21.5	21.6	11.1
NW	25	25.2	5.7
HX	16.1	17.3	18.1
HS	16.4	22	6.3
LL	29.8	29.9	27.3
NT	7.4	13.7	49.6
TB	23.7	23.8	8.9
VG	19.1	19.2	16.2
BB	30.3	30.4	16.9
DB	28.9	28.9	28.9
Mean	*20.6*	*22.6*	*22.1*

Comparing the results in Tables 3 and 4 , it is difficult to settle on any general characteristics that might distinguish the two result sets. In general, the Butterworth filter attributes more noise to the overall signal, although they are of the same order. The implication is that it is sufficient to use a simple filtering method (Moving Average) rather than a more complex method such as Butterworth. The complexities of the latter tend to be embedded in software procedures, whereas code

for the Moving Average method is easily visible. The results in Table 4 also show significant variation between organisations. This variation simply reflect differing volatility in the original reputation signal. Overall, the S/N values are higher than for the Kalman filter, but are low in comparison to S/N values encountered in good audio signals (approximately 70–100 db). Further comments are in Section 5.

4.3. Special Treatment of Power Spectra Arising from Reputational Shocks

Two organisations, VOL and HS in Tables 3 and 4, suffered one severe reputational shock each in the period 2014–2015. They were both involved in illegal activities which were widely reported in the press at the time, and there was much adverse comment on social media. Long-lasting reputational downturns resulted. The power spectra of the unfiltered signals show apparently exceptionally high noise levels in the order of 70%. This does not seem reasonable, and we think that the shock levels are so severe that all non-shock parts of the signal appear to be "noise". To give a more reasonable estimate of the noise level, the outer quartiles of the signal (i.e., the lowest 25% and the highest 25%) of the scores were excluded from the power spectrum assessment. This process resulted in noise extimates of approximately 80%.

4.4. Monetary Impact

In Mitic (2017b), several results for the monetary effect of reputation are given. We consider, as an illustration, the figures given for monetary gain or loss in stressed circumstances. This is not a "worst case". It represents either very negative or very positive reputation, but not extreme cases of either, which are very rare. Suppose that, in general, positive reputation accounts for P_+% of profit after tax, and that negative reputation accounts for P_-% of profit after tax. Suppose, further, that the *mean maximum* and *maximum maximum* percentage noise levels for filter F considered are $M_{F,mean}$ and $M_{F,max}$ respectively. F takes the values K (for Kalman), B (for Butterworth) or MA (for Moving Average). Then, we attribute percentage monetary amounts with the products $(P_+ \ M_{F,mean}/100)$ and $(P_+ \ M_{F,max}/100)$ as the mean and maximum noise associated with positive reputation respectively over a one year period. The expressions for negative reputation are similar: P_+ is replaced by P_-. In addition to the parameter F for the filter method, let parameter α refer to the "mean/max" designation (so it takes values *mean* or *max*), and let parameter β refer to "positive/negative" reputation (so it takes values "+" or "−"). Then, Equation (24) gives an expression for $V_{F,\alpha,\beta}$, the revised percentage monetary value of the reputation component of profit after tax when noise is removed.

$$V_{F,\alpha,\beta} = M_{F,\alpha} \ (1 - \frac{P_\beta}{100}) \tag{24}$$

The numerical values of the variables used in this section are:

1. $M_{K,max} = 9.9\%$ (from Section 4.1)
2. $M_{K,mean} = 7.6\%$ (from Section 4.1)
3. $M_{B,max} = 22.6\%$ (from Section 4.2.2)
4. $M_{B,mean} = 20.6\%$ (from Section 4.2.2)
5. $M_{A,max} = 30.7\%$ (from Section 4.2.1)
6. $M_{A,mean} = 10.9\%$ (from Section 4.2.1)
7. $P_+ = 2.3\%$ (from Mitic (2017b): Table 5 average)
8. $P_- = 2.7\%$ (from Mitic (2017b): Table 5 average)

Table 5 gives the values obtained by applying Equation (24) with the numerical values above. Values in the table are given correct to 1 decimal place, which is an appropriate accuracy in this context.

A few points are notable from Table 5. First, there is very little differentiation between the results for positive and negative reputation, and for the "mean" and "maximum" estimates. That differentiation is much more marked in cases of extreme reputational stress, for which the value

of the parameters P_+ and P_- are of the order 5% and 10% respectively (see Mitic 2017b, sct. 4.2). Even so, the absolute monetary amount can be substantial. Suppose, for example, that profit after tax is about $500 m. Using the "maximum" Kalman figures from Table 5, positive reputation can add 2.1%, or $10.5 m, to that profit. Negative reputation can remove 2.5%, or $12.5 m from it.

Table 5. Revised Reputation Contributions to Profit after Tax using mean and Maximum Noise Estimates.

Filter	Reputation	Max % Noise Estimate	Mean % Noise Estimate
Kalman	Positive	2.1	2.1
Kalman	Negative	2.5	2.4
Butterworth	Positive	1.8	1.8
Butterworth	Negative	2.1	2.1
Moving Average	Positive	2	1.6
Moving Average	Negative	2.4	1.9

5. Discussion

The aim of this study was to measure the extent of noise in a reputational time series. The Kalman filter is particularly suited for this purpose because it is possible to adjust a filtering parameter specifically for assessing measurement noise, and because its formulation admits an appropriate simple static model. The limiting nature of the degree of smoothing applied allows us to identify the noise component in the observed signal. The results shown in Table 2 indicate that a figure of about 10% would be a reasonable estimate for a maximum noise level applicable to all data sets considered. The mean noise level, a rounded value for which is 8%, is only slightly less. These estimates arise from setting the critical value for the confidence interval in Equation (15) to 99%. If the 95% critical value is used instead the noise percentages may be reduced by approximately one quarter. The S/N values in Table 2 are extremely low compared to what would be expected of a good audio signal (70–100 decibels). Listening to a reputation signal would be like listening to a whisper in a lot of hiss! The Kalman filter results are compared with those derived from low pass filters, principally using the Butterworth filter. Filters other than Kalman do not really assess the amount of measurement noise: they measure total noise. Consequently they are not as suitable as the Kalman filter in answering the question "How much noise does the measurement process add?".

Overall, the result is useful because it enables us to identify a neighbourhood of the neutral value of the observed signal that is unreliable. Sentiment in this region conveys very little reputational information, and may be interpreted as "chatter". These small deviations from the neutral value can therefore be discarded in further analyses.

Conflicts of Interest: The author declares no conflict of interest.

Appendix A. Calculation of System Noise Parameter q

Using the standard deviation of the observed signal as an estimate of the process noise in the Kalman analysis may be justified by noting the stability of the volatility of the signal, measured on overlapping time slices. For each organisation, the standard deviation of the reputation scores in 100-day (approximately three months) periods, overlapping by 50 days, were noted. The mean of those standard deviations was then compared with the standard deviation of the scores taken over the entire time. The results are noted in Table A1. They indicate that, to an acceptable degree of accuracy, that the standard deviations are stable over time. Therefore the standard deviation taken over the entire time is a reasonable estimate of process noise.

Table A1. Comparison of Mean of Time-sliced Standard Deviations and overall Standard Deviation of Reputation Signals.

Organisation	Mean Time-Sliced SD	Overall SD
BW	0.049	0.048
FT	0.078	0.076
FD	0.058	0.058
HUD	0.075	0.072
NSS	0.066	0.067
MB	0.049	0.047
VOL	0.097	0.094
BSG	0.115	0.112
RYB	0.112	0.108
NW	0.108	0.099
HX	0.098	0.091
HS	0.117	0.099
LL	0.143	0.135
NT	0.124	0.114
TB	0.150	0.147
VG	0.141	0.137
BB	0.087	0.085
DB	0.093	0.084

Appendix B. Kalman Calculation Parameters

Table A2 shows calculated Kalman parameters for each organisation. In that table, column *Relative % Error* is the relative error $100(y_t - x_t)/y_t$, where y_t and x_t are defined in Section 3.2. The Kalman calculations were implemented in Excel, and the maximum likelihood calculations used to calculate the Box-Cox parameter λ were implemented in R.

Table A2. Kalman Calculation Parameter Values.

Organisation	Relative % Error	Parameter q Estimate	Parameter r Estimate
BM	1.11	0.05	7.3
FT	0.11	0.08	4.31
FD	0.82	0.06	6.53
HUD	1.26	0.07	16.04
NSS	0.56	0.07	6.18
MB	0.65	0.05	8.58
VOL	0.29	0.1	4.36
BSG	0.39	0.11	10.5
RYB	0.17	0.11	5.79
NW	0.26	0.11	5.19
HX	0.29	0.1	9.83
HS	0.16	0.12	6.65
LL	0.35	0.14	9.9
NT	0.2	0.12	8.08
TB	0.26	0.15	7.69
VG	0.25	0.14	6.22
BB	0.28	0.09	9.19
DB	0.48	0.09	4.66

References

Butterworth, Stephen. 1930. On the Theory of Filter Amplifiers. *Wireless Engineer* 7: 536–41.

Cerna, Michael, and Audrey F. Harvey. 2009. The Fundamentals of FFT-Based Signal Analysis and Measurement, National Instruments Application Note 041. Available online: http://www.ni.com/white-paper/4278/en/ (accessed on 8 November 2017).

Cosma, Ioana A., and Ludger Evers. 2010. *Markov Chains and Monte Carlo Methods: Lecture Notes, Chapter 7*. African Institute for Mathematical Sciences (AIMS). Available online: http://www.isn.ucsd.edu/classes/beng260/complab/week2/Cosma_Evers_2010.pdf (accessed on 6 February 2018).

Deloitte RiskAdvisory. 2016. Reputation Matters: Developing Reputational Resilience Ahead of Your Crisis. Available online: https://www2.deloitte.com/content/dam/Deloitte/uk/Documents/risk/deloitte-uk-reputation-matters-june-2016.pdf (accessed on 8 November 2017).

Hamilton, James. 1986. State-space models. In *Handbook of Econometrics, Chapter 50*. Amsterdam: Elsevier, vol. 4, pp. 3039–80.

Hamming, Richard Wesley. 1989. *Digital Filters*. Englewood Cliffs: Prentice-Hall. Available online: https://www.dsprelated.com/freebooks/filters/ (accessed on 6 February 2018).

Jackson, Leland B. 1996. *Digital Filters and Signal Processing: With MATLAB Exercises*. Boston: Kluwer Academic Publishers.

Jung, Jae C., and Su Bin Park. 2016. Case Study: Volkswagen's Diesel Emissions Scandal. *Thunderbird International Business Review* 59: 127–37. Available online: http://onlinelibrary.wiley.com/doi/10.1002/tie.21876/full.

Kalman, Rudolph Emil. 1960. A New Approach to Linear Filtering and Prediction Problems. *Transactions of the ASME–Journal of Basic Engineering* 82: 35–45.

Liu, Bing. 2015. *Sentiment Analysis: Mining Opinions, Sentiments and Emotions*. Cambridge: Cambridge University Press.

Mitic, Peter. 2017a. Standardised Reputation Measurement. In *Proceedings of the IDEAL 2017, Guilin, China*. Edited by Hujun Yin, Yang Gao, Songcan Chen, Yimin Wen, Guoyong Cai, Tianlong Gu, Junping Du, Antonio J. Tallón-Ballesteros and Minling Zhang. Lecture Notes in Computer Science (LNCS). Cham: Springer, vol. 10585, pp. 534–542. doi:10.1007/978-3-319-68935-7_58.

Mitic, Peter. 2017b. Reputation Risk: Measured. In Proceedings of the Paper present at the Complex Systems May 2017, The New Forest, UK, 23–25 May 2017. Edited by George Rzevski and Carlos Brebbia. Southampton: Wessex Institute of Technology Press.

Oppenheim, Alan V., and Ronald W. Schafer. 1989. *Discrete-Time Signal Processing*. Upper Saddle River: Prentice Hall.

Shumway, Robert H., and David S. Stoffer. 2017. *Time Series Analysis and Its Applications: With R Examples*, 4th ed. New York: Springer International, chp. 6.

Smith, Julius Orion. 2007. *Introduction to Digital Filters: With Audio Applications*. Berlin: W3K Publishing.

Smith, Steven W. 1999. *The Scientist and Engineer's Guide to Digital Signal Processing*, 2nd ed. Poway: California Technical Publishing. Available online: http://www.analog.com/en/education/education-library/scientist_engineers_guide.html (accessed on 6 February 2018).

Stoica, Petre, and Randolph L. Moses. 2005. *Spectral Analysis of Signals*. Upper Saddle River: Prentice Hall.