Reputation: Probability Distributions, Prediction and Simulation

Peter Mitic

***Abstract*—The economic significance of reputation in the context of a proposed distribution of reputation scores is discussed. Proposals are made to use the distributional properties of reputation for prediction and simulation. A method of expressing reputation numerically is presented as a weighted average of sentiment scores derived from multiple contents within a given time window. Given a sufficiently extensive reputation time series, averaging induces a marked clustering near to a modal value. The proposed *bi-exponential* distribution models this clustering better than other candidate distributions. The economic effects of a specific reputational shock is examined to illustrate both its severity, persistence and subsequent consequences.**

***Index Terms*—Reputation, probability distribution, simulation, goodness-of-fit, biexponential.**

I. Introduction

The concept of *reputation* is often used in a loose sense - that of the dictionary definition. The Cambridge English Dictionary (*https://dictionary.cambridge.org/dictionary/english/reputation*) defines reputation as “the opinion that people in general have about someone or something...”. However, that definition is too imprecise to permit formal measurement. The purpose of this paper is to use a rigorous quantitative definition of reputation to formulate a reputation distribution. That distribution can be used to make predictions and run simulations to study the possible effects on an organisation of reputational events. More generally, sub-structures within a reputation distribution can shed light on the general effect of reputation on the economic and business relationship of an organisation with its stakeholders.

We therefore start with a brief discussion of the economic importance of reputation, and then place informal concepts of reputation and sentiment within a formal quantitative context. Candidate reputation distributions are then assessed. The results show that a proposed *bi-exponential* distribution passes goodness-of-fit tests adequately. Indications are then given of how that distribution can be used to analyse the effect of reputational events.

II. Literature Reviews

This literature review is divided into two parts: one for reputation as it relates to economics, and the other for reputation measurement.

Manuscript received December 11, 2021; revised May 3, 2022.

P. Mitic is with the Dept. of Computer Science, University College London, Gower Street, London WC1E 6BT (e-mail: p.mitic@ucl.ac.uk).

A. Reputation and Economics

Firms were urged in the 1990s to consider reputation as a part of normal business practice. At that time reputation was defined loosely in the dictionary sense, and was measured by survey. The advantages of maintaining a good reputation were, and still are, expressed ultimately in monetary terms, albeit not quantified. A summary of early work is given by Fombrun [1], in which the conceptual history of reputation is traced. The point is made that consumers rely on firms’ reputations because they have nothing else to inform them on products and services. A firm, on the other hand, can use its (good) reputation in marketing. The economic benefits that are claimed to ensue from a good reputation are summarised by Cannon [2]. They mainly amount to reduced costs. Increased sales is not included, but is clearly relevant. Caminiti [3] provides a further point. A company with a good reputation is seen as being altruistic with respect to society. Cole [4] attempted to value reputation by sourcing reputation scores from survey data, and linking them to company balance sheet items (e.g. EBIT, EPS, Yields etc.). This technique predates direct measurement by automated content collection, and is somewhat subjective. Furthermore, balance sheet items are subject to many other factors (supply, labour, markets etc.).

The YouGov survey and report (Rowe [5]) is a good summary of more recent qualitative views: 76% of respondents agreed that reputation is linked to the overall financial performance of their organisations. More recently, Valenzuela [6] identifies three alternative factors that have influenced reputation: globalisation, sustainability and the digital revolution. He argues that these factors constitute a ’new normal’ as a means of doing business, and that there is a consequent social contract between organisations and their stakeholders. Finally, he identifies a means of measuring reputation which was not possible prior to widespread use of the internet: direct procurement of content by data mining, followed by Natural Language Processing (*NLP*).

B. Reputation Measurement

Reputation measurement relies on analysis of opinion and sentiment. In particular, *NLP* has become a significant component of that analysis. An early start was made by Droba [7], who worked on measuring public opinion in the 1930s. Droba used questionnaires containing selected statements that expressed sentiment. Respondents were asked to rank them in order of sentiment by pairwise comparisons. Ratings were then attached to the ranks using an arbitrary scale.

The review by Donsbach and Traugott [8] documents progress in opinion measurement from the 1940s to the up to the end of the 20th century. Opinion polls, conducted by interview and later by telephone, were pioneered by George Gallup in 1930 [9], and grew in extent in the 1940s and 1950s.

They sought opinion mostly on political and consumer issues. Early polls were criticised for not using representative samples, and for not distinguishing between fact and opinion. Small advances were made in sampling techniques (such as Gallup's 'open' and 'closed' question mix), but the main contribution made throughout the second half of the 20th century was the amount of data collected. A mid-term critique of polling in that period may be found in Margolis [10], who questions whether or not respondents are sufficiently informed in their answers, and whether or not they tell the truth. Brooker, and Schaefer [11] provide an update on polling methods in the 1990s, including early use of the internet, probability sampling, bias, randomness and sampling error.

Basic categorisation of sentiment as *positive*, *negative* or *neutral* was achieved by the start of the 21st century, in parallel with increasingly widespread use of the internet. A comprehensive review of the development of sentiment analysis may be found in Mantylaa [12] or Liu [13]. Liu gives more specific details and examples. Dave [14] gives details of an early application in the context of product reviews. That context remains a major target for sentiment analysis today. By 2015, support for a continuous sentiment metric had been developed by identifying and quantifying emotions expressed in text (Cambria [15]).

The techniques used for sentiment analysis have also progressed markedly as it became possible to procure extensive amounts of data. The principal methods employed are summarised in Godsay [16], and Jurafsky [17] provides specific methodological details. In particular, Turney's (unsupervised) *Pointwise Mutual Information, Information Retrieval* (*PMR-IR*) method (Turney [18]) is a well-established unsupervised method. It accounts for the semantic and syntactic context of text, as well as word frequency. The use of *Naive Bayes* analyses, documented in Lewis [19], has produced robust supervised *Hidden Markov* models that use bigrams, despite assumptions of event independence. Machine Learning techniques have gained ground in the past decade (see Boiy [20] for example). Although some good results can be obtained, semantic analysis by machine learning is difficult because of factors such as language ambiguity, unstructured text, slang and implied meaning.

III. Reputation and Sentiment Measurement

Measurement of sentiment and reputation by direct collection of textual contents was explained in Mitic [21]. A summary is given here. The technique depends on exploiting *Natural Language Processing* (*NLP*) methods to elucidate and quantify the sentiment expressed in each content. For a discussion of *NLP* techniques, see Liu [13] and Jurafsky [17]. The results for multiple contents are then combined to produce a *local reputation score*. Fig. 1 sets out the major steps.

Thus, for an entity G, and a single content c received in a time period t, denote the sentiment expressed in that content by a real number, $s(c, G, t)$ in the interval [-1,1]. $s(c, G, t) > 0$ represents positive sentiment and $s(c, G, t) < 0$ represents negative sentiment. Nominally, t is one day.

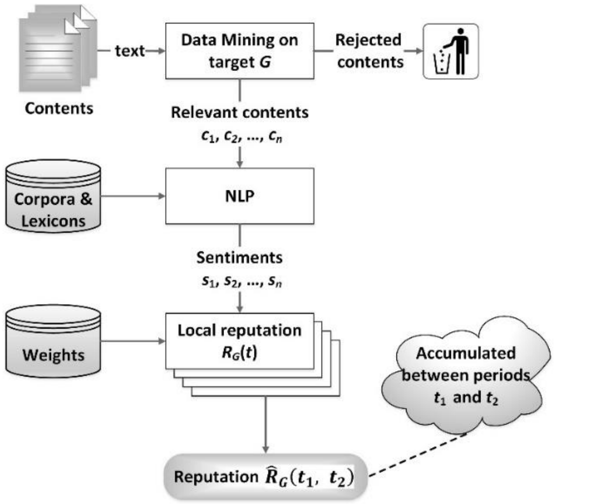

Fig. 1. Overall process for calculating sentiments and reputation.

The *local reputation* of G by the end of t, $R_G(t)$, follows as a weighted average of sentiments. The first step is to collect a set of n contents $\{c_i\}$ $\{i = 1,2....,n\}$ with corresponding sentiments $\{s_i(c_i, G, t)\}$ $(i = 1,2....,n)$, all for the same target in the same time period. The next step is to define weights $\{w_i\}$ $(i = 1,2....,n)$, dependent on the influence of the content sources and the mode of transmission of those contents. Then the *local reputation* at time t is given by Equation 1.

$$R_G(t) = \frac{\sum_{i=1}^{n} w_i s_i(c_i, G, t)}{\sum_{i=1}^{n} w_i} \quad (1)$$

Long term *reputation*, $\hat{R}_G(t_1, t_2)$, is then an extensive (at least 6 months is recommended) time series of values of R_G for values of t between t_1 and t_2.

$$\hat{R}_G(t_1, t_2) = \{R_G(t) : t = t_1, t_{1+1}, \ldots, t_2\} \quad (1a)$$

If the time period $t_1..t_2$ is extensive, we would drop the t-arguments, and refer to $\hat{R}_G$ to indicate a 'long term' reputation. The analysis that follows is based on values of $R_G(t)$.

A. Modal Concentration of Reputation: Observations

We have observed that distributions of local reputation scores are highly concentrated about a central point which approximates to the modal score. Attempts to reduce distributions of period reputation scores to Normal distributions using a Box-Cox transformation are successful in some cases, but not all. Even 'successful' cases which pass a goodness-of-fit test for normality, display a density surrounding the mode that is more than would be expected from a Normal distribution. Informally, reputation distributions resemble exponential mixtures. Section *V*(*B*) shows an example.

There is evidence of the same type of concentration from other sources. Tran [22] notes convergence of sentiment in a social network to a 'group sentiment' value as more people contribute to the network. Given an issue under discussion, as time increases, more people comment on the issue, and a consensus sentiment with respect to the issue emerges. Average sentiment for the group is seen to converge to a limit which represents the sentiment for the group as a whole. The forms of the traces in Fig. 2 and Fig. 4 of Tran [22] appear to be consistent with exponential convergence. It is likely that 'group sentiment' is a key idea in explaining the shape of reputation distribution.

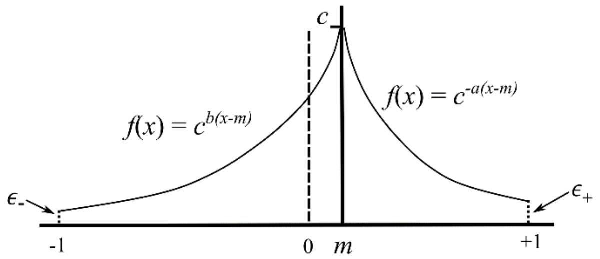

Fig. 2. *BiExponential* density.

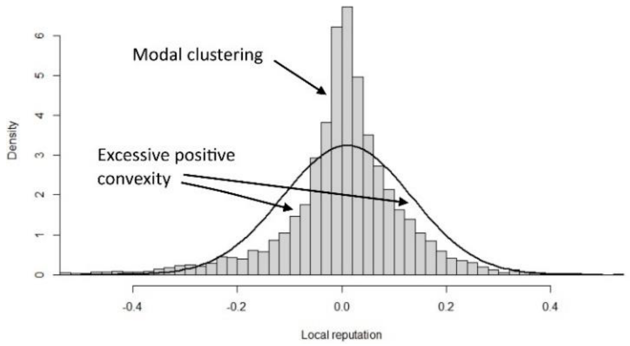

Fig. 3. Empirical distribution of local reputation scores (all data sets combined), with fitted Normal density.

The same profiles of sentiment plotted against time were noted in Reagan [23]. Fig. 6 in that paper shows positive and negative sentiment scores separately. Both exhibit exponential-like convergence. Frequency is shown using a log scale for, leading to an extremely high sentiment score concentration near a point thar represents neutral point. Outliers are rare and are sparsely distributed.

A less severe case is reported the study of three million *tweets* by Mozetic *et al* [24]. Approximately two thirds of tweets were rated neutral in human trials, with approximately one sixth rated positive and one sixth rated negative (Fig. 10 in the Mozetic paper). There is no indication of a more detailed sentiment distribution, but the percentage of neutral sentiments is significantly large.

IV. Reputation Distribution

In this section we suggest potential distributions for the *local reputation* score, $R_G(t)$. In Section *V*(*B*) we settle on one of them on the basis of best fits to data. We first note the pattern in which reputational activity on issues emerges and decays. That provides motivation for a distributional model using exponential distributions.

A. Modal Concentration of Reputation: Explanation

The phenomenon of modal clustering can be explained by partitioning sentiments into two sets. The first represents "background" sentiment: non-extreme and usually centered on a value near zero. The second represents prominent issues. Particular issues often invite much more extreme sentiment (either positive or more often negative), resulting in a much more extreme *local reputation* score. Some emerge quickly and gain rapid traction, whereas others emerge slowly and grow gradually. As time advances, some issues have a long period of stability of extreme sentiment, and others peak and decline almost instantly. For both, a decay phase follows as interest wains, and the decay rate varies from issue to issue. As a general rule, issues gain traction rapidly and decay at a slower rate, but there are exceptions. The simulations in the *Results* section shows an example of how a reputation time series can be simulated by superimposing a small number of extreme sentiments on a large volume of low-level ”background” sentiment.

In order to address the issues described above, we propose three candidate distributions. Each is bipartite so that reputation scores greater than the modal value can be modelled separately from reputation scores less than the modal value. The principal contenders are the *Bi-Exponential* and the *HyperNormal* distributions. They are discussed first. Both are compared with a bipartite half normal distribution (hereinafter referred to as *HalfNormal*), which is intended to be an improvement on a Normal distribution. A Normal distribution for this reputation data is known to be a poor model near the modal point, as illustrated in the *Results* section.

B. Reputation Density: Bi-Exponential

The remarks in the preceding section prompt a model of a reputation time series using a pair of exponentials - the *bi-exponential* distribution. The distribution comprises an exponential distribution with rate $b > 0$ for reputation values x less than or equal to a reference value m, and a second exponential distribution with rate $a > 0$ for reputation values x greater than or equal to m. The parameter m is typically near zero and can be positive, negative, or zero. Figure 2 shows the form of the density, $f(x)$. The distribution is truncated at $x = \pm 1$. In all cases observed so far, the values of a and b are such $f(-1)$ and $f(1)$ are both marginally positive. That is, two very small values ϵ_+ and ϵ_- can be found with $f(\epsilon_+) > 0$ and $f(\epsilon_-) > 0$. In practice, ϵ_+ and ϵ_- are both approximately 10^{-5}.

The equation of the *BiExponential* density is (with $m \in (-1,1)$, $a > 0$ and $b > 0$, and a normalising constant c),

$$f(x,m,a,b) = \begin{cases} c\ e^{b(x-m)} & ; x \in (-1, m] \\ c\ e^{-a(x-m)} & ; x \in [\,m, 1) \end{cases} \tag{2}$$

The normalising constant, c, is given by

$$c = \frac{a\ b\ e^{b(1+m)}}{-a + a\ e^{b(1+m)} + b\ e^{b(1+m)} - b\ e^{b(1+m)+a(m-1)}} \tag{3}$$

When fitting the density to data, two constraints $f(1,m,a,b) = \epsilon_+$ and $f(-1,m,a,b) = \epsilon_-$ are added to force the fitted curve to pass sufficiently close to the points $(1, \epsilon_+)$ (1,+) and $(-1, \epsilon_-)$ respectively.

The *BiExponential* distribution function, $F(x) = \int_{-\infty}^{x} f(t)dt$, is (with m, c, a and b as above):

$$F(x,m,a,b) = \begin{cases} \frac{c}{b}\left(e^{b(x-m)} - e^{-b(1+m)}\right); & x \in (-1, m] \\ \frac{c}{b}\left(1 - e^{-b(1+m)}\right) + \frac{c}{a}\left(1 - e^{-a(x-m)}\right); & x \in [\,m, 1) \end{cases} \tag{4}$$

C. Reputation Density: HyperNormal

The proposed *HyperNormal* distribution is a variant of the Normal distribution in which the term x^2 is replaced by a more general x^n ($n > 0$). That replacement, with an appropriate value for n, can produce a density that is inflated for small values of x and is deflated for larger values of x. That serves as a better model for the observed modal peak. The degree of inflation near $x = 0$ is also controlled by an additional parameter, s. In most cases $n \in [0,1]$ achieves the desired degree of inflation, but in a few cases a value of n marginally greater than 1 provides a better fit to data. The HyperNormal density and distribution functions are given in Equations 5 and 6 respectively. Both are defined for $x \geq 0$ only, so negative and positive reputation scores have to be modelled separately, relative to the modal value m, as with the *BiExponential* distribution.

$$f(x,s,n) = \frac{1}{s\,2^{\frac{1}{n}}\Gamma\left(1+\frac{1}{n}\right)}\, e^{-\frac{x^n}{2\,s^n}};\ x \geq 0; s,n > 0 \tag{5}$$

$$F(x,s,n) = 1 - \frac{1}{\Gamma\left(\frac{1}{n}\right)}\,\Gamma\left(\frac{1}{n},\frac{x^n}{2\,s^n}\right); x \geq 0; s,n > 0 \tag{6}$$

(The 2-parameter Gamma function in Equation 6 is the upper incomplete gamma function $\Gamma(w,x) = \int_x^\infty t^{w-1}e^t dt$)

Although the absolute empirical reputation scores are in [0,1] whereas the *HyperNormal* density extends to infinity, the probability that $x > 1$ is so small for the values of s encountered that they can be ignored.

D. Reputation Density: HalfNormal

Empirical studies (see the *Results* section) show that the Normal distribution is a poor fit for reputation scores. The Normal distribution is not an adequate model for either the modal peak, nor for distribution asymmetry. A possible alternative is a bipartite distribution comprising two half-normal distributions. The *HalfNormal* distribution is defined for positive values of the independent variable only, so negative and positive reputation scores have to be modelled separately relative to the modal value m (as with the previous proposals). The *HalfNormal* density and distribution functions are given in Equations 7 and 8 respectively.

$$f(x,\sigma) = \frac{\sqrt{2}}{\sigma\sqrt{\pi}}\, e^{-\frac{x^2}{2\,\sigma^2}};\ \ x \geq 0;\ \sigma > 0. \tag{7}$$

$$F(x,\sigma) = erf\left(\frac{x}{\sigma\sqrt{2}}\right)\ ;\ \ x \geq 0; s > 0. \tag{8}$$

The same comments on the region $x > 1$ that were made for the *HyperNormal* distribution also apply for the *HalfNormal* distribution.

E. Goodness-of-Fit Tests

We propose two *GoF* tests. The first is the *TNA*-test, described in Mitic [25], which is a formalisation of a *QQ*-plot, uses the raw reputation data directly, and was specifically developed for cases where rare events with significant impacts are possible. This test is independent of the sample size used, and is therefore a good alternative to tests such as *Kolmogorov-Smirnov*, which often provide no discrimination for sample sizes of more than 100 (i.e. all cases that we considered). The *TNA*-test is unusual in that the test statistic is a direct measure of *p*-value, and that a low test statistic value indicates a good fit (value zero is a perfect fit).

For comparison, a second *GoF* test, based on the *Student t* statistic, is used. Data are partitioned by size, and a prediction of the mean datum is made for each partition. Each prediction is paired with the corresponding empirical mean value within its partition. A paired value *t*-test can then be used to test *GoF*. The results depend on the number of partitions, so this *t*-test is less satisfactory than the *TNA* test. Despite this restriction, the two generally agree.

V. Results

A. Data

The data are daily reputation metrics (the values $R(G,t)$ defined in Equation 1 with the period t set to 1 day) sourced from the business intelligence organization *alva* (*alvagroup.com*). *Alva* provides objective data that covers thousands of UK corporates since 2014. The *BiExponential* distribution has been applied to a representative selection of 19 corporates, using sufficient data to cover a 6 month period at least. The data are normalised to the range [-1,1], positive values corresponding to overall positive sentiment in the day concerned, and negative values corresponding to negative sentiment.

B. Empirical Distribution

The illustration in Fig. 3 shows the empirical density for all 19 data sets used for the analysis in this paper. It illustrates two points relative to a fitted Normal distribution (both indicated on the illustration). First, there is a marked clustering near to the modal peak. Second, the empirical distribution has an increased convexity. The distribution shown is typical of distributions for individual data sets.

C. GoF Results

Table I shows the *TNA* p-value (described previously) and paired value *t*-test results when fitting the *BiExponential* and *HyperNormal* distributions. Fits using a *HalfNormal* distribution are shown for comparison. Organisations are referred to anonymously in column *Org*. In the *TNA* columns, p-values **less than** 0.05 indicate an acceptable fit (i.e. with at least 95% confidence). Values less than the 0.01 are 'excellent' fits. In the *t* columns, p-values **greater than** 0.05 indicate an acceptable fit. In both cases *GoF* fails at 95% confidence are indicated in bold typeface.

TABLE I: *TNA* AND *t*-TEST P-VALUES FOR THE *BIEXPONENTIAL*, *HYPER-NORMAL* AND *HALFNORMAL* FITS

	BiExponential		*HyperNormal*		*HalfNormal*	
Org	*TNA*	*t*	*TNA*	*t*	*TNA*	*t*
A1	0.039	0.132	0.024	0.350	0.021	0.060
A2	0.030	0.202	0.027	0.240	0.032	0.107
B1	0.012	0.208	0.021	0.312	**0.054**	**0.013**
B2	0.010	0.243	0.028	0.218	**0.068**	0.070
B3	0.025	0.185	**0.050**	0.197	**0.092**	**0.001**
B4	0.018	0.385	0.031	0.308	**0.067**	0.054
B5	0.024	0.422	0.044	0.370	**0.088**	0.051
B6	0.037	0.237	0.048	0.217	**0.066**	0.064
B7	0.015	0.229	0.038	0.497	**0.093**	**0.003**
B8	0.031	0.228	0.047	0.115	**0.085**	0.069
B9	0.014	0.365	0.035	0.409	**0.069**	**0.002**
B10	0.012	0.379	0.023	0.470	**0.058**	**0.036**
M1	0.014	0.210	0.016	0.291	0.044	0.071
M2	0.029	0.119	0.041	0.298	**0.086**	**0.002**
M3	0.040	0.118	**0.065**	0.243	**0.110**	**0.005**
M4	0.013	0.173	0.035	0.340	**0.077**	**0.010**
M5	0.018	0.269	0.031	0.303	**0.082**	0.067
M6	0.022	0.456	0.035	0.401	**0.077**	0.108
L1	0.013	0.166	0.024	0.393	**0.069**	0.184

The results in Table I indicate that the *BiExponential* and *HyperNormal* distributions outperform the *HalfNormal* distributions. The *HalfNormal* distribution fails *GoF* tests for too many examples.

Fig. 4 shows a comparison of distributions fitted to a typical organisation: B2 in Table I. The empirical data density (derived from a histogram) is shown in black, and the three candidate densities are superimposed in the shades indicated. The modal value is shown by the vertical dashed line. The *HalfNormal* distribution fails to capture the peak in the neighbourhood of the modal value, and, more generally, tends to have a concave rather than a convex curvature for mid-size reputation values. The *HyperNormal* distribution has the same curvature problem, although there is an improvement in capturing the modal peak. The *BiExponential* distribution correctly models the curvature and is much more successful at modelling the modal peak. For comparison, a Normal distribution fit is also shown. For that Normal distribution to be at all reasonable, the calculated (using maximum likelihood) standard deviation was halved. The Normal fit in Fig. 3 does not use such an amended standard deviation.

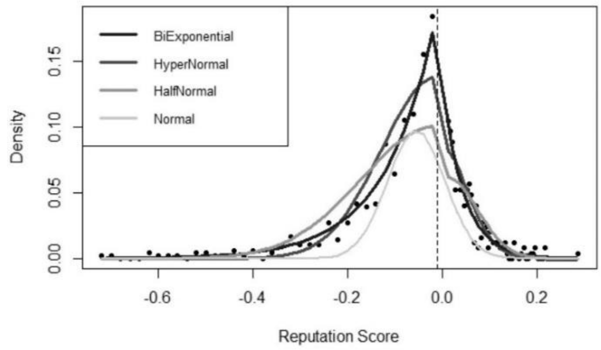

Fig. 4. *BiExponential*, *HyperNormal*, *HalfNormal* and Normal fits for organisation B2.

D. Simulations

The distributional properties of a reputation score can be used to effectively to simulate data. Even if the nominal measurement period is extensive, it may be beneficial to have more for statistical analysis. Simulated data can offer exactly that, provided that the simulated data are a valid proxy for the actual data. In the discussion that follows we give a brief indication of how a simulated reputation distribution can be used. There is a significant difference in the analysis of significant negative and significant positive reputational events. In both cases, the *BiExponential* distribution is an optimal base for simulation.

1) Negative shocks

Additional distributional properties are needed to investigate the effects of significant negative reputational events. Opinion holders and agents of transmission (the press, social media etc.) react to reputational events in different ways, sometimes generating a 'reputation shock' due to massive negative sentiment. The shock is expressed as an extreme and rapid change in *local reputation*. Very few extreme shocks have been observed to date. Notable instances are the Volkswagen 'Dieselgate' scandal (September 2015) and the Boeing 737-Max air crashes (October 2018 and March 2019). The major features of such shocks are listed below.

- The intensity of the shock (i.e. the maximum absolute value of the local reputation at shock inception)
- The time from inception to peak intensity
- The time that peak intensity persists
- The reputation profile and relaxation time from peak intensity to an ambient pre-shock level
- The frequency and profile of after-shocks

Reputation profiles differ in the way the above features are expressed. Some issues emerge rapidly and there is an immediate shock. Others are slower. The period at peak intensity generally lasts between 1 and 14 days. After that the relaxation period is usually much longer, and is a slow reversion to an ambient reputation level. Fig. 5 shows an example: the 'Dieselgate' profile for Volkswagen (M3 in Table I). The simulated *BiExponential* shock profile is very similar to the actual. It shows the initial rapid 3-day reputation drop at day 80, a 15-day 'low' period, a slow 90-day exponential recovery, with periodic interruptions by after-shocks with annual frequency 10. The consequences were dire. Volkswagen's net profits in 2014, 2015, 2016 and 2017 in m€ were respectively 12697, -4069, 7103 and 13818 (see https://www.statista.com/statistics/272053/operating-profit-of-volkswagen-since-2006/). A three-year recovery period is apparent.

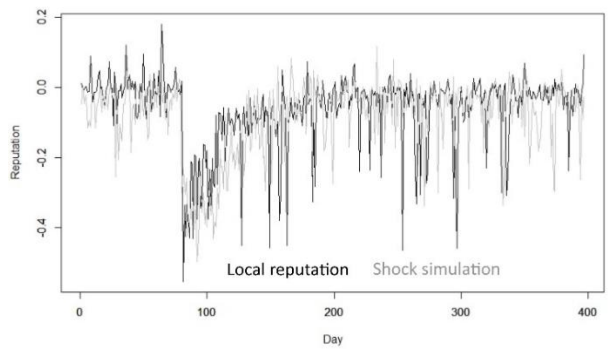

Fig. 5. Volkswagen: original reputation data (black), *BiExponential* shock simulation (grey).

2) Positive Shocks

Positive reputational shocks follow a completely different pattern. They tend to occur as short-lived bursts, lasting only a few days or even one day only. Consequently, the effect on long-term reputation is negligible compared to the effect of a single large negative reputational shock. An appropriate model is therefore to set a peak value, generate short runs positive sentiment values, and apply them at random intervals. Fig. 6 shows an example based on Mercedes-Benz (M4 in Table I).

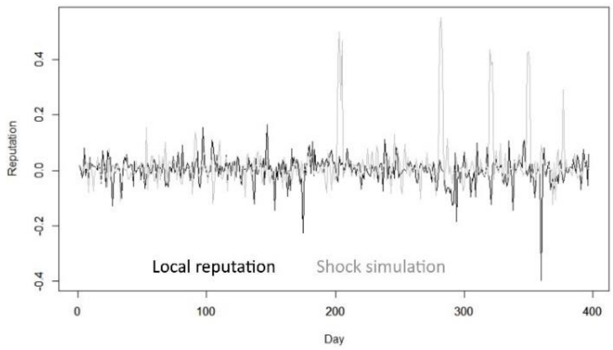

Fig. 6. Organisation M4: original data (green), *BiExponential* simulation (grey) and positive shock applied to simulation (black).

VI. Discussion: Wider Reputational Issues

A. Consequences of Poor Reputation

Some consequences of poor reputation are apparent. A recent extreme case is that of Cambridge Analytica (*CA*). In 2018 *CA* harvested the personal data of millions of Facebook profiles without consent and used it for political advertising. This was contrary to UK data protection law. *CA* was forced to cease trading because of a massive loss of confidence in their actions, and trust in social media providers in general dropped (see *https://www.businessinsider.com/facebook-twitter-socialmedia-confidence-charts-20188?r=US&IR=T*). Facebook's share price fell for six months after the scandal was uncovered.

A similar, but less drastic case emerged in April 2020. The insurer Hiscox refused to pay claims related to Covid-19 Business Insurance (see https://www.insurancetimes.co.uk /news/hiscox-faces-legal-action-over-rejecting-coronavirus-biclaims/1433137.article). Their share price halved from a high point in January 2020 (when they were regarded as trustworthy) to the beginning of May 2020.

The Volkswagen (*VW*) 'Dieselgate' affair in September 2015 is an example of an almost instant loss of reputation (as measured by direct procurement) from an ambient level of zero (neutral reputation) to a low point -0.55 (-1 is the worst possible). An account of its development may be found in Contag [26]. A provision for remediation of 6.5bn EUR was made (later increased to 16.2), and an approximate calculation in Mitic [27] shows that *VW* lost approximately £55m directly in lost sales. After about 4 months the reputation level recovered near to its ambient level, but with periodic downturns as old news resurfaced. The *Results* section of this paper shows an illustration.

B. The Effect of Social Media

A notable development in the decade 2010-2020 is the rise of social media as a means to communicate and to influence. Suarez [28] has a report on the effect of social media on corporate reputation in which reputation values are sourced from the Spanish reputation index *MERCO* (*Monitor Empresarial* de *Reputacion Corporativa*). This index is survey-based, so content coverage is necessary limited, and the results represent a single snapshot in time. The authors found that the total number of news items in social media has a small positive effect on corporate reputation. The study did not investigate the effect of negative items, although the examples in the previous section show that negative content can be disastrous. The study in Mitic [29] shows that the effect of negative content (not wholly based on social media) has a much more serious effect than positive content. However, the point is also made that reputation can be used to bolster a firm's reputation. Factors such as customer engagement, cost-benefit analysis and the impact of social media communication are discussed in Floreddua [30].

Press articles often give a more direct view of the effect of social media. The news article *https://www.bbc.co.uk/news/business-48871456* discusses factors such as hacking, fake news, rumour, influence of the originator, scams, cyber-crime, and handling negative content. The effect of consumer organisations that undertake product reviews remains anecdotal. They appear to influence consumers, but any effect has not yet been measured. UK examples of such organisations include *The Consumers' Association* (a registered charity, informally known as *Which?*), the retail money management website *moneysupermarket.com* and BBC Radio's money management programme *Moneybox.* Similarly, the precise effects of review websites such as *TripAdvisor*, are yet to be explored.

VII. Conclusion

Reputational time series have the peculiar property of being very tightly clustered near a model value. This makes it difficult to fit a uni-partite probability distribution. A good fit near the model value tends to compromise the fit elsewhere, and vice versa. We have demonstrated that the proposed bipartite *BiExponential* distribution is usually an optimal fit. Using the *BiExponential* distribution, it is possible to simulate time series that closely resemble the originals from which they were derived. This approach allows further investigation of the results of particular reputational events. In particular, severe negative reputational shocks can have devastating consequences for the finances of an organisation. It is therefore of value, when deciding on major policy changes, to investigate the likely reputational effects of those changes. Examples include introducing a new product to the market, or associations with other organisations.

Severe negative reputational effects are often long lasting. In contrast, severe positive reputational effects tend to have little lasting effect other than to help maintain an overall positive mean reputation. It is therefore more important to attempt to avoid management decisions that might lead to negative sentiment, than to rely on actions that could generate positive sentiment.

Acknowledgment

The author is grateful for the support of alva-Group for their continued interest, support, and assistance in the preparation of this paper.

References

[1] C. J. Fombrun and C. van Riel, "The reputational landscape," *Corporate Reputation Review*, vol. 1, no. 1, pp. 5-14, 1998.

[2] H. M. Cannon and M. Schwaiger, "The role of company reputation in business simulations," *Simulation and Gaming*, vol. 36, 188-202, 2005.

[3] S. Caminiti, "The payoff from a good reputation," *Fortune*, vol. 125, no. 3, pp. 49-53, 1992.

[4] S. Cole, "The impact of reputation on market value," *World Economics*, vol. 13, no. 3, 2012.

[5] O. Rowe. (2014). YouGov presentation to the Public Relations Consultants Association. [Online]. https://yougov.co.uk/topics/consumer/articles-reports/2014/11/12/economics-reputation

[6] A. L. Valenzuela, 2018, *The Connecting Leader*, Lioncrest.

[7] D. D. Droba, "Methods used for measuring public opinion," *American Journal of Sociology*, vol. 37, pp. 410–423, 1931.

[8] W. Donsbach and M. W. Traugott, *The SAGE Handbook of Public Opinion Research*, SAGE Publications London, 2007.

[9] G. Gallup, "A scientific method for determining reader-interest," *Journalism Quarterly*, vol. 7, no. 1, pp. 1-13, 1930.

[10] M. Margolis, "Public opinion, polling, and political behavior," *Annals of the American Academy of Political and Social Science*, vol. *472*, pp. 61-71, 1984.

[11] R. Brooker and T. Schaefer, "Public Opinion in the 21st century: Let the people speak?" *Cengage Learning*, Boston MA, 2005.

[12] M. V. Mantylaa, D. Graziotin, and M. Kuutila, “The evolution of sentiment analysis,” *Computer Science Review*, vol. 27, pp. 16-32, 2018.

[13] B. Liu, “Sentiment analysis: Mining opinions,” *Sentiments and Emotions*, CUP Cambridge UK, 2015.

[14] K. Dave, S. Lawrence, and D. M. Pennock, “Mining the peanut gallery: Opinion extraction and semantic classification of product reviews,” in *Proc. 12th int. Conference on the WWW*, 2003, pp. 519–528.

[15] E. Cambria, P. Gastaldo, F. Bisio, and R. Zunino, “An ELM-based model for affective analogical reasoning,” *Neurocomputing*, vol. 149, pp. 443-455, 2015.

[16] M. Godsay, “The process of sentiment analysis: A study,” *Int. Jnl. Computer Applications*, vol. 126, no. 7, pp. 26-30, 2015.

[17] D. Jurafsky and J. H. Martin, “*Speech and Language Processing*, Pearson London, 2008.

[18] P. D. Turney, “Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews,” in *Proc. 40th ACL Meeting Philadelphia,* 2002, pp. 417-424.

[19] D. D. Lewis, “Naive (Bayes) at forty: The independence assumption in information retrieval,” *ECML-1998, LNCS 1398,* Springer, 1998.

[20] E. Boiy and M. Moens, “A machine learning approach to sentiment analysis in multilingual Web texts,” *Inf Retrieval*, vol. 12, pp. 526–558, 2009.

[21] P. Mitic, “Can reputation replace regulation?” in *Proc. ICEFR 2020 Paris*, 2020, vol. 11, no. 5.

[22] H. Tran and M. Shcherbakov, “Detection and prediction of users attitude based on real-time and batch sentiment analysis of facebook comments,” *Computational Social Networks*, pp. 273-284, 2016.

[23] A. J. Reagan, C. M. Danforth *et al.*, “Sentiment analysis methods for understanding large-scale texts: a case for using continuum-scored words and word shift graphs,” *EPJ Data Sci*., vol. 6, no. 28, 2017.

[24] I. Mozetic, M. Grcar, and J. Smailovic, “Multilingual Twitter sentiment classification: The role of human annotators,” *PLoS ONE*, vol. 11, no. 5, 2016.

[25] P. Mitic, “Improved goodness-of-fit tests for operational risk,” J. *Operational Risk*, vol. 15, no. 1, pp. 77-126, 2015.

[26] M. Contag, G. Li *et al.*, 2017, “How they did it: An analysis of emission defeat devices in modern automobiles,” in *Proc. IEEE Symposium on Security and Privacy (SP), San Jose, CA*, 2017, pp. 231-250

[27] P. Mitic, “The effect of conduct risk losses on reputation,” *Conduct Risk – the Definitive Guide, ed. Peter Haines*, Incisive Media London, chapter 10, 2016.

[28] L. M. Suarez, J. P. Lopez, and B. C. Saiz, “The influence of heuristic judgments in social media on corporate reputation,” *Sustainability*, vol. 12, no. 4, 2020.

[29] P. Mitic, “Standardised Reputation Measurement,” in *Proc. IDEAL 2017*, Guilin China, 2017, pp. 1–9

[30] P. B. Floreddua, F. Cabiddua, and R. Evaristo, “Inside your social media ring: How to optimize online corporate reputation,” *Business Horizons*, vol. 57, no. 6, pp. 737-745, 2014.

[31] I. Mozetic, M. Grcar, and J. Smailovic, “Multilingual Twitter sentiment classification: The role of human annotators,” *PLoS ONE*, vol. 11, no. 5, 2016.

Peter Mitic is Honorary Professor in the Department of Computer Science, University College London, and is Head of Operational Risk at Santander Bank UK. He has many years of experience in financial risk analysis and has a particular interest in the quantification of reputation, its statistical properties, and in how it can influence economic and business decisions.