



Research Paper

Reputation risk contagion

Peter Mitic

Santander UK, 2 Triton Square, Regent's Place, London NW1 3AN, UK;
email: peter.mitic@santander.co.uk

(Received October 7, 2016; revised December 8, 2016; accepted January 3, 2017)

ABSTRACT

The effects of the reputation of any single member of a group of agents on all the others in the group are calculated by modeling the spread of reputation contagion in a DeGroot network. The reputation of individual agents is measured by compiling a reputation index for each agent over an extended period. Transition probabilities within the network are assessed by considering extreme reputational events using a Bayesian approach. The results indicate that consensus is reached quickly, and influential agents can be easily identified. Agents in the network with a very positive reputation serve to mitigate the negative reputation of other agents in the network. Approximately 10–15% of the reputation of any agent in the network is attributable to network effects; positive reputations are deflated and negative reputations are inflated. The network effect on the sales of any single agent can be estimated once the reputation score has been translated to sales.

Keywords: reputation risk; reputation risk measurement; correlation; DeGroot network; Bayesian; consensus.

1 INTRODUCTION AND MOTIVATION

Previous studies (Perry and de Fontnouvelle 2005; Fiordelisi *et al* 2013, 2014) measured the effect of reputational events for corporate organizations by linking those

events with share price changes. They used the term “reputation risk” to indicate that share capital will suffer as a result of an adverse reputational event. In a study covering the years 2014–15, Mitic (2017) extended the work of these authors, and established three important principles: reputation can be measured objectively; reputation can be positive as well as negative; and there is a correlation between corporate sales and reputation. The last point is important because it provides the motivation to maintain a good reputation. A positive reputation enhances sales, whereas a negative reputation depresses sales, the latter being a more marked effect than the former. The aim of the present study is to assess the effects of the reputation of the members of a group on any single member of the group. There are many ways in which the word “effects” in the previous sentence may be interpreted, as the following questions demonstrate.

- (1) Can the reputation of any one organization be boosted/diminished by other organizations?
- (2) How fast do reputational effects spread within a group or organization?
- (3) How much of the reputation of one organization is attributable to other organizations?
- (4) Which organizations in the group are the most (and least) influential?

We aim to answer all of the above questions in this analysis.

The motivation for a network analysis comes first from observed correlations in the performance and behavior of particular organizations, and second from informal comments that may or may not affect decisions made by customers. The latter is hard to measure in isolation, but the task is made easier by a network analysis. The bases of the DeGroot model are the ideas of “learning” and “social influence”, from which a contagion effect may be inferred. It is therefore an appropriate model for quantifying contingent actions between agents in a group.

1.1 Aims and structure of the paper

This paper presents an approach to network contagion analysis that is novel for two reasons. First, a means to analyze contagion is reputation. A way to measure reputation has only recently been developed, and there have been few opportunities to explore applications of this measure. Second, the reputation measure is used to elucidate a network structure. Nothing is assumed about the structure, or the weights that it represents, a priori. In contrast, other studies (see Section 2.5) assume that the network structure is known in advance. The particular point stressed in this paper is that consensus is reached on reputation via contagion, and the result is that the proportion of reputation due to systemic factors can be estimated.

Section 2 gives a brief review of research on the essential elements – reputation, contagion, influence and consensus – with an explanation of the contribution each makes. In particular, the DeGroot consensus model is introduced.

In Section 3, the methodology used to measure reputation is introduced, and the DeGroot method is explained in detail. Section 4 gives an account of the empirical results obtained by applying the DeGroot analysis to the reputation scores of a network of ten UK retail banks. Section 5 revisits a previous result (Mitic 2017) for quantification of reputation, and explains how those results should be amended by the effects of the network. General conclusions are drawn in Section 6.

1.2 Terminology and notation

In the theoretical and general discussions that follow, we will use the term “agent”, since this is more consistent with the terminology used in network theory. In other cases, terms such as “bank” or “organization” will be used, particularly if those terms make sense in context.

1.3 Reputation, contagion, influence and consensus

It is assumed that agents within a social network interact and influence one another to a greater or lesser extent. In this context, “network” should be interpreted in a very general sense, such that its agents include individual people, information sources (television, internet, the press, etc) and corporates. Within that categorization, we focus on the relationship between banks and their customers. Customers formulate opinions about their banks, and the banking sector in general, by “customer–bank” and “customer–other agent” interactions. Such interactions are subject to influence and opinion, thereby spreading contagion in the form of the opinions of influential actors in the network. The term “contagion” usually has negative connotations (disease, rumor, etc), but in this context it should simply be interpreted as the spreading of a sentiment of any kind, positive as well as negative. Such sentiment can be measured (the subject of Section 3.1), and this measurement is expressed as “reputation”.

The key elements in this analysis are therefore linked as follows.

- (1) Actors within a network influence each other by canvassing opinions from, and listening to, agents they are connected to in the network. They then combine the opinions of network neighbors with their own opinions.
- (2) Contagion, in its broadest sense, is spread through such interactions.
- (3) Opinions can be detected and measured, from which the concept of reputation arises.

- (4) The degree of influence of any one agent on any other can be measured in terms of the reputation measure by looking at parallel reputational movements for distinct agents. There is an assumption that significant movements in the reputation of a pair of agents implies that one has influenced the other.
- (5) Consensus may be achieved in the network, such that we can consider the network has its own opinion on an individual actor within it. The DeGroot model is a well-established mechanism for measuring that consensus in social networks.

An extensive discussion of interactions in social networks may be found in Bruder *et al* (2014). In particular, Bruder *et al* identify two principal pathways to consensus: mimicry and evaluation of signals provided by others in the network. The concept of reputation is missing from their analysis.

2 LITERATURE REVIEW

This section contains a review of research that focuses on three of the main points discussed in the previous section: contagion, influence and consensus. The subject of reputation, and its measurement in particular, is a relatively new area of study, and will be discussed in Section 3.

2.1 Contagion effects in banking networks

Some previous studies of contagion effects applied to banking have adopted DeGroot or similar methodologies. Most have used alternative methods, largely dependent on formulating a network in advance of further calculations. This section gives a brief summary of recent work.

Modeling contagion by considering a sequence of credit defaults is a common theme. Examples are Battiston *et al* (2013), Gai and Kapadia (2010), Beale *et al* (2011) and Elliott *et al* (2014). These studies all share common features. First, their networks are determined by obligor-counterparty relationships, and are therefore pre-determined. In contrast, the network in our study is implied by reputational interaction calculations. Second, the principal aim in these studies is to estimate a measure of systemic risk. An interesting result due to Gai and Kapadia is that when the degree of network connectivity is 6 or more the increase in contagion is negligible. This is a significant result, and agrees with our results in Section 4.

A completely different technique is used by Aldasoro *et al* (2015), who analyze contagion due to illiquid assets and external influences by formulating an optimization problem, where liquidity and external monetary influences are the origin of contagion. Once again, the emphasis is on debt default.

Further studies deal with other aspects of bank interactions, and all use either a highly connected or a fully connected network. Hałaj and Kok (2014) and Degryse and Nguye (2007) model interbank lending and borrowing; Chan-Lau *et al* (2007) model shock propagation through the network. Finally, Tabak and de Castro Miranda (2013) use a partitioned network to isolate certain nodes.

2.2 What causes contagion in banking?

What causes contagion is beyond the main scope of this paper, but some pointers are given in this section. The emphasis in all cases is on shock propagation and systemic risk.

The Organisation for Economic Co-operation and Development (OECD 2012) attributes contagion to responses to cross-border financial shocks. Similarly, Brown *et al* (2014) attribute contagion to withdrawals at one bank triggering withdrawals at another bank. In a study aimed more at isolating the effects of systemic risk, Dungeya and Gajurela (2015) conclude that crisis shock transmission is responsible for systemic contagion.

2.3 The DeGroot method for determining consensus

The DeGroot (1974) model is a network analysis of social interaction, information transmission and consensus formation. It is represented by a matrix of weights, T , and opinion formation is a process on a graph. The matrix T is an input to a Markov chain process, whereby they may be updated. A clear summary of the method, with simple examples, may be found in Jackson (2008).

The idea behind the DeGroot process is given by DeMarzo *et al* (2003), although it is expressed in terms of “listening” rather than influence. In a multistage process (in terms of time), each agent listens to the opinions of all other agents. They start at time $t = 0$ with an initial opinion vector $p(0)$, which describes their opinion on some particular topic. At $t = 1$ they listen to the opinions of their neighbors, and assess the precision of those opinions. As a result they may or may not update their own opinions, thereby deriving a new opinion vector $p(1)$. The process is repeated at $t = 2$, resulting in the revised opinion vector $p(2)$, and so on. A similar idea is expressed by Golub and Jackson (2010), who model a situation in which, at each point in time, the agents’ beliefs are updated by taking a weighted average of the beliefs of their neighbors. This model is applicable for large networks in which highly influential agents are not prominent.

In a DeGroot process, the opinion vector is updated by considering neighbors of nodes that are not directly connected to the node. Rather than listen, agents look, and are influenced by what they see. Pan (2012) notes that a DeGroot learning mechanism

implies full accessibility of information: each agent is able to observe all other agents' opinions.

Pan (2010) also provides an account of the DeGroot method, in which agents continuously update their opinions by taking weighted averages of those of the neighboring agents. Pan's varying influence weight matrix T is therefore time dependent, but the base DeGroot model is the same. The context of the influence model described in Pan and Gilles (2010) is an extension of the model outlined in this section. Their model is cast in terms of game theory, with a cost of update at every stage, and an overall strategy for each agent.

2.4 DeGroot network analysis in other contexts

Before considering how the DeGroot method is applied in the context of reputational risk, a brief review of applications in other fields is merited. In particular, the technique has been used primarily in the context of social networks, and often with a twist.

Friedkin and Johnsen (1990) extended the DeGroot model to account for an agent's own opinion being distinct from the opinion of other agents. Their conclusions are largely as expected: an agent is influenced most when the opinions of all other agents differ from its own, and the outcome for an individual agent is determined by the difference between the total effect of other agents and its own effect.

Das *et al* (2014) apply a mixture of the DeGroot and Flocking models. In the latter, agents assign more weight to opinions that already conform to their beliefs. They also use the Voter model, in which an agent chooses a neighbor's opinion at random and moves to this opinion. This is an empirical study based on responses to surveys, designed to find a "best-fit" network model.

Ghaderi and Srikant (2012) use the Perron–Frobenius theorem explicitly to model stubborn agents in a network. They set up a cost function (with a stubbornness parameter) that has to be minimized with respect to the sum of squares of differences between the opinions of distinct agents. Their overall result was that the opinions of non-stubborn agents converge to a linear combination of the opinions of the stubborn agents, provided those agents are stubborn enough.

An extension of the DeGroot model due to Friedkin is used by Jia *et al* (2015) to model the situation in which an individual agent values its own belief over and above the influence of others. Their method is very similar to that used in this paper as an extension to the principal analysis. We also consider other cases. The precise methodology is given in Section 4. The Jia *et al* (2015) model, perhaps not surprisingly, predicts convergence to the view of the dominant agent.

2.5 Reputation measurement

The starting point for the measurement of reputation is tokenization, followed by part-of-speech (POS) tagging and then sentiment assessment. Tokenization is the process of categorizing easily identifiable parts of a sentence (eg, the symbols \$, ?) and splitting a whole text into sentences, and sentences into words. POS tagging associates each word with a part of speech. POS tagging is particularly important in identifying adjectives and adverbs, which tend to convey the sentiment expressed in the text. The process starts by finding possible parts of speech using a corpus: an extensive list of sentences in which each word has a POS pre-attached. The groundwork was laid by Kučera and Francis (1967), who formulated the Brown corpus. Since 1967 many others have been developed, notably the Penn Treebank (Marcus *et al* 1993) and the Corpus of Contemporary American English (COCA; Davies 2010). A lookup in a corpus is not sufficient to fully identify the POS of every word. For example, the word “good” could be a noun or an adjective. To resolve this difficulty, a common technique is to find the most common usage in the corpus. For example, in the Brown corpus “good” occurs 61 times as a noun and 693 times as an adjective. A simple lookup would therefore classify “good” as an adjective every time. In more sophisticated analyses, neighboring words can be used to provide a context. Viterbi’s algorithm (Viterbi 1967; Forney 1973) is still widely used for this purpose.

The next step in measuring reputation is to assess the sentiment expressed, first by analysis of the words in each sentence and then by aggregating the sentiments on all sentences in a text. Liu (2015) gives a full summary of research on sentiment analysis. In addition, Pang *et al* (2002) discuss identification of the holder (who expresses the sentiment, usually a noun), and the target (the aim of the sentiment, also usually a noun). The principal task in sentiment analysis is to determine the “polarity” of individual sentences: positive, negative or neutral. A supplementary task is then to assess any emphasis within a sentence that might serve to amend an initial assessment of polarity. Good summaries are given by Jurafsky and Martin (2008) and Liu (2015), who discuss the concept of “nearness” to key words to determine sentiment. Commonly used alternative techniques used to determine sentiment are naive Bayes, maximum entropy and support vector machines (Chaudhari and Govilkar 2015).

Reputation is generally considered to be a measure involving an accumulation of sentiment scores, usually with corresponding weights, but there is no academic consensus on how the calculation should be done. Collomb *et al* (2014) use a weighted average of sentiment scores, with weights that reflect the importance of each sentiment score. Peleja *et al* (2014) use weights read from a lexicon (ie, a dedicated corpus). The Alva method described in Section 3.1.2 uses weights that reflect the influence of the source of the text under consideration.

2.6 Other measures of consensus in banking

In this section, we give a brief review of alternative methods for determining network consensus. They reinforce the principal concepts used in this paper: reputation, contagion, influence and consensus.

Some variations on the DeGroot method have already been mentioned (Jackson 2008; Golub and Jackson 2010; Friedkin 2011). Commenting about the paper by DeMarzo *et al* (2003), the above-mentioned authors incorporate the concept that an agent is influenced not only by immediate influential neighbors, but also by influential neighbors within a given radius (in network terms) of the agent.

Jackson and Lopez-Pintado (2013) model influence and contagion with a diffusion process, governed by a diffusion differential equation. Each agent is assigned a state, and the way in which an agent is influenced by other agents is determined by how many neighbors are in the same state. The diffusion aspect sets our study apart from others, and amplifies the concept of contagion spread, which is central to our paper.

The idea of the state of an agent is taken much further by Grabisch and Rusinowska (2011), in a largely theoretical study of influence based on aggregation functions. They examine the interplay between an agent's own opinion (termed "inclination") with the opinion of others, and reinforce the concepts of influence and convergence in a network in a novel way by defining and evaluating several aggregation functions.

Several other concepts used in the analysis in Section 3.3 have been employed in other contexts. Georg (2014) established a foundation for Bayesian analyses in the context of a model of lines of credit in the interbank market. A threshold similar to that introduced in Section 3.3.3 was employed by Londono (2016) to measure contagion as the proportion of international stock markets that simultaneously experience unexpected returns beyond a certain limit. Finally, as an alternative to the DeGroot convergence mechanism, Herrera-Viedma *et al* (2002) present a consensus model with a formal feedback loop in the context of multi-person decision making.

3 METHODOLOGY: REPUTATION AND INFLUENCE

This section sets out the theoretical basis for determining influence in a social network, and describes the practical considerations in reputation measurement.

3.1 Reputation measurement

The measure of reputation risk used in this analysis is the index produced by the business intelligence company Alva (see www.alva-group.com). A discussion of their methodology may be found in Mitic (2017), and a summary of it follows. More generally, considerable advances in sentiment analysis have been made in the past

twenty years, and a good overview and summary may be found, for example, in Liu (2015).

3.1.1 Reputation risk definitions

In order to set the context for what may be described colloquially as “reputation risk”, we use the general idea that stakeholders have expectations of an organization, and sometimes the organizations meet those expectations and sometimes they do not. There is therefore a potential gap between performance and expectation. This gap is mostly interpreted as “performance falls short of expectation” (hence the term “risk”), but it is quite clear that performance can exceed expectation. If performance exceeds expectation, “risk” is something of a misnomer, but we can continue to regard good performance as “risk”, perhaps calling it “negative risk”. The following definitions therefore suffice informally.

Reputation: a perception of an organization on the part of stakeholders that can affect, positively or negatively, the business relationship between the stakeholder and the organization.

Reputation event: an occurrence or action that affects reputation.

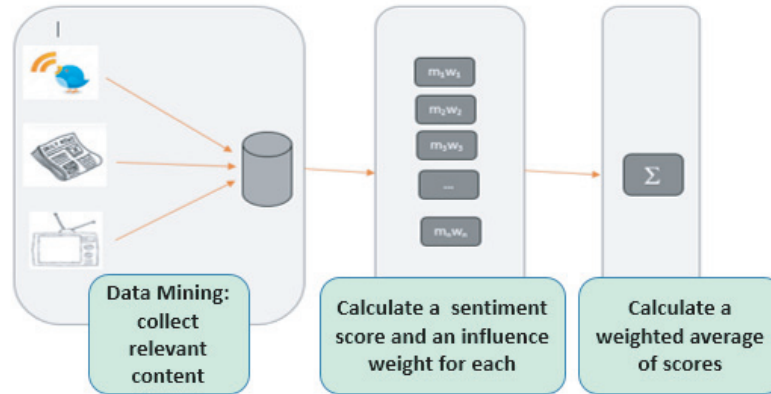
Reputation risk: the difference between stakeholder expectation and organization performance.

Reputation risk measurement: numerical assessment of reputation.

The definition of “reputation risk” given above is similar to that provided by Federal Reserve Bank of Boston (1995), except that the latter does not account for positive reputation.

3.1.2 Reputation risk measurement

The Alva Sentiment Index is compiled daily, and provides a short-term measure of “sentiment” (used interchangeably with “reputation” in this paper). Given target words or phrases that are relevant to an organization i , live electronic feeds supply public data from news media, social media, people or organizations whose opinion carries weight, as well as trade reports and surveys. This stage is termed “content harvesting”. Each item received (termed “content”) within a twenty-four-hour period (“day t ”) is scored on a scale of 1, . . . , 10 for four factors: overall sentiment, influence of the source, prominence of the organization and relevance. A weight is assigned to each score, representing factors such as the importance of the source of the content, the prominence of the organization within the content and the extent to which the positive or negative nature of the content is stressed. A weighted mean of the scores

FIGURE 1 Stages in index construction: content harvesting, analytics, index compilation.

with respect to the four factors is then calculated. The result of this procedure is a score m_{jit} , between 1 and 10, for organization i relating to a content j on day t . Further weights, reflecting the importance of the content as a whole, are then assigned to each content. National media or influential persons are weighted most highly. For organization i and content j , let the weight on day t be w_{jit} . The final index value $aRI_i(t)$ (for Alva Reputation Index) is a weighted average of the scores for all contents for organization i received on day t :

$$aRI_i(t) = \sum_j m_{jit} w_{jit}.$$

A numerical result of 5.5 (the median of 1 and 10) is a “neutral” score. An index value greater than 5.5 represents “positive” sentiment, whereas an index value of less than 5.5 represents “negative” sentiment. The index is reset to 5.5 at the start of each twenty-four-hour period. The value of the index is typically between 4.8 and 6.2, and values outside that range are rare. Figure 1 shows a schematic of the construction process. In Figure 1,

- (1) the “content harvesting” (ie, data mining) block represents an accumulation of content from the media, with database storage;
- (2) the “content analytics” block shows calculation of the weighted scores $m_{jit} w_{jit}$ for content i ; and
- (3) the “index” block represents the aRI construction stage.

The aRI is particularly suitable for the network analysis that follows because the sentiment measure is condensed into a single number, and also because it provides a fast-moving view of reputation. The RepTrak index, compiled by the Reputation Institute (see www.reputationinstitute.com), is less suitable because it comprises a suite of measures, and cannot be compiled quickly as it is survey based. The construction and validation of the RepTrak index is described in Ponzi *et al* (2011) and Fombrun *et al* (2015).

In the quantitative network model that follows, the data comprises a daily aRI value for each organization concerned.

3.2 Network analysis

The essential concepts in network analysis are summarized in Appendix A online. Here, the DeGroot influence model is introduced. A significant difference between theoretical presentations of network analysis and reality is that the network for the actors in this study is, from the beginning, unknown.

3.3 The DeGroot method: details

Details of the DeGroot consensus model are given in the first section of Appendix B online. Here we present a summary of the equations involved, with brief notes on the notation used throughout the paper.

The basis of the DeGroot method is an influence matrix, T , which represents the extent to which each agent in a network can influence all the agents in the network, including itself. The entry T_{ij} in the i th column and j th row of T represents the weight that agent i places on the current belief of agent j in forming agent i 's opinion ($0 \leq T_{ij} \leq 1$). An opinion vector $p(t)$ expresses the a priori opinions of agents in the network at a time t . Equation (3.1) gives the DeGroot recurrence relation for $p(t)$, and (3.2) gives $p(t)$ explicitly in terms of the initial opinion vector $p(0)$. The limiting value of $p(t)$ for large t , if it exists, is denoted by $p(\infty)$ (see (3.3)):

$$p(t) = Tp(t-1), \quad (3.1)$$

$$p(t) = T^t p(0), \quad t = 1, 2, \dots, \quad (3.2)$$

$$p(\infty) = \lim_{t \rightarrow \infty} (T^t) p(0) = T_\infty p(0). \quad (3.3)$$

3.3.1 Convergence in a DeGroot network

The second section of Appendix B online contains a discussion on convergence in a DeGroot network. For consensus to be reached, the condition that must be satisfied by the opinion vector $p(\infty)$ is $Tp(\infty) = p(\infty)$. Convergence implies a consensus belief, derived from the agents' initial beliefs. In this paper, necessary and sufficient

conditions for convergence are not strictly needed, as it is quick and easy to demonstrate convergence empirically. Appendix B.2 (available online) notes that a useful indicator that consensus has been reached is that there exists some value of $t (> 0)$ such that all entries in the columns of T^t are the same. This is explored in Section 4.

3.3.2 The Friedkin “amendment”

Friedkin (2011) proposed a model for stubbornness: the influence of any individual agent is treated as superior to all other agents. Appendix B.3 online provides a brief explanation. Equation (3.1) is replaced by

$$p(t) = DTp(t-1) + (I - D)p(0), \quad (3.4)$$

which contains a “stubbornness” matrix D . The expression $p(\infty)$ in (3.3) is replaced by

$$p(\infty) = (I - AT)^{-1}(I - D)p(0). \quad (3.5)$$

The results using this approach are given in Section 4.5.

3.3.3 Influence measured by sentiment movement

In order to apply the DeGroot theory in the context of aRI values, we introduce the concept of “sentiment movement” to measure influence. Such movement is calculated as the daily differences in sentiment scores. In the following discussion it should be noted that the use of the term “agent” in this context is actually more complex than it might first appear. There is, first of all, an organization at which comments are directed, and those comments are detected and analyzed as described in Section 2.2. The commentators collectively constitute the “agent”. It is convenient to refer to the commentators, and also to the organization upon which they comment, as the “agent”. With this dual reference in mind, the influence of agent i on agent j is built as follows.

If the value of the aRI for agent i on day t is $\text{aRI}_i(t)$, then the movement from the previous day, $M(i, t)$, is given by

$$M(i, t) = \text{aRI}_i(t) - \text{aRI}_i(t-1). \quad (3.6)$$

Large differences in movement are notable because they indicate changes in sentiment that originate from a large number of sources. Smaller differences are considered to be “noise”. The aRI value is not often moved by any single item (new report, blog, etc), even if that item originates from an influential source. Therefore, for each agent i , we identify all movements greater than or equal to a “high” threshold L_H . In addition, we also identify all movements greater than or equal to a “very high” threshold L_{VH} , which is greater than L_H .

Therefore, let the count of movements, $C(i, \lambda)$, for agent i for days $1, \dots, n$, subject to a limit λ , be defined by

$$C(i, \lambda) = \{M(i, t) : \text{abs}(M(i, t)) \geq \lambda, 1 \leq t \leq n\}, \quad \text{where } \lambda = L_H, L_{VH}. \quad (3.7)$$

Setting numerical values for L_H and L_{VH} will be discussed in the next section. The statistical justification for setting the values for L_H and L_{VH} is that these limits should represent “milestones” in statistical significance, namely the 5% and 1% significance limits usually applied in significance tests. Therefore, we designate 5% of the largest absolute empirical movements as “large”, and calculate L_H accordingly. Similarly, we designate 1% of the largest absolute empirical movements as “very large”, and calculate L_{VH} accordingly. If the 5% and 1% figures are varied, the conclusions in the following sections still stand, although the numerical values in (3.10) are not the same. An important invariant is the structure of the networks in Figures 5 and 6, which is unchanged, subject to the constraints discussed in Section 4.6, as L_H and L_{VH} vary.

Influence is measured by first choosing a particular agent (agent i) and counting the number of large movements, and then pairing agent i with a different agent (agent j), and counting the number of corresponding very large movements in the sentiment of the agent j , given that there was a large movement in the sentiment of agent i . Dependency is measured by a simple conditional probability. Denoting the influence of agent i on agent j ($i \neq j$) by T_{ij} , we have

$$\begin{aligned} T_{ij} &= C(j, L_{VH}) \mid C(i, L_H) \\ &= \frac{C(j, L_{VH}) \& C(i, L_H)}{C(i, L_H)}. \end{aligned} \quad (3.8)$$

The idea behind the expression in (3.8) is that a large movement in the sentiment of agent i , associated with a very large movement in the sentiment of agent j , implies that agent i has influenced agent j . In terms of the commentators referred to at the start of this subsection, “people who have commented on organization i have influenced people who have commented on organization j ”. The entry T_{ij} in (3.8) is calculated by counting the large movements for agent i on day t , locating all movements for agent j on the same day and then counting how many of the latter are very large. T_{ij} is the count of those very large movements for agent j divided by the count of the large movements for agent i . An alternative formulation for T_{ij} would condition the influence of agent i on agent j on past movements of reputation of agent j .

When $i = j$, the above discussion does not apply in the same way. It must be remembered that “agent” in this context means a collection of commentators, any one of which can influence any other. However, the same ratio, as described above, can be used with a different interpretation. Given two counts $C(i, L_H)$ and $C(i, L_{VH})$,

the right-hand side of (3.8) reduces to

$$T_{ii} = \frac{C(i, L_{VH})}{C(i, L_H)}. \quad (3.9)$$

This ratio gives a measure of the influence of the commentators responsible for a very large sentiment movement on the commentators responsible for a merely large sentiment movement. The case $i = j$ should be regarded as a measure of the extent to which agent i values its own opinion.

When there are n agents in total, the collection of T_{ij} values for $1 \leq i, j \leq n$ constitutes the $n \times n$ influence matrix T .

The differentiation between a large movement for the sentiment of agent i and a very large movement for the sentiment of agent j in (3.8) highlights what is assumed to be a significant influence of agent i on agent j . In other words, agent i has a very large effect on agent j . An alternative view, in which a very large movement in the sentiment of agent i influences a large movement in the sentiment of agent j , would be a much weaker consideration.

3.4 Monetary gain/loss due to reputation

This section contains a summary of the justification for the view that “reputation means money” and then develops the theory of total influence, whereby reputation contagion effects propagated through the network are related to monetary losses or gains.

3.5 The link between reputation and sales

A methodology for determining the link between the reputation measure aRI and annual sales was described in Mitic (2017). The aim of that study was to demonstrate that an organization can lose sales if it suffers from a poor reputation, but can also gain sales if it enjoys a good reputation. A very brief summary is given here, and a summary of results may be found in Section 5. These results are modified following the DeGroot analysis.

With the assumption that reputation can affect future sales, correlations were sought of aRI movements with lagged changes in sales, the lags being between zero and six weeks. Having established the existence of such correlations separately for negative and positive sentiment, best fit curves were found such that the change in sales resulting from a change in sentiment could simply be read from the fitted curve. In most cases a lognormal curve was the best fit. To model cases where sentiment is extreme, random samples were drawn from the distributions represented by the best fit curves, and 99.9% value-at-risk was calculated. “Super stressed” sentiment cases were also defined in terms of scenarios. These represent cases where reputation (and thus the aRI) is either sustainably positive or sustainably negative.

3.5.1 Total influence

The first stage in calculating the effect of network sentiment on sales is to formulate a long-term view of an agent's reputation using a cumulative excess sentiment ($CES(j)$) measure for an agent j . $CES(j)$ is simply an accumulation over the times (nominally days) $t = 1, \dots, n$ of the differences between an actual sentiment measure $aRI_j(t)$ (see Section 3.1.2), and a nominal "par" value of the sentiment measure, \overline{aRI}_j , which is independent of time. Formally,

$$CES(j) = \sum_{t=1}^n (aRI_j(t) - \overline{aRI}_j). \quad (3.10)$$

For a set of M agents indexed by $j = 1, \dots, M$, define $\boldsymbol{\Gamma}$, a vector of CES values:

$$\boldsymbol{\Gamma} = \{CES(j)\}_{j=1}^M. \quad (3.11)$$

$\boldsymbol{\Gamma}$ is an objective measure of the perceived reputation of the M agents. The long-term influence matrix \boldsymbol{T}_∞ (see (3.3)) provides a long-term consensus view of sentiment. It shows the relative influence of one agent on another. \boldsymbol{T}_∞ has columns in which all elements are the same (to a given accuracy). Let \boldsymbol{J} be a vector whose entries are the column values of \boldsymbol{T}_∞ . We now define the total influence of the system over y years, τ , by the scalar product

$$\tau = \frac{\boldsymbol{\Gamma} \cdot \boldsymbol{J}}{y}. \quad (3.12)$$

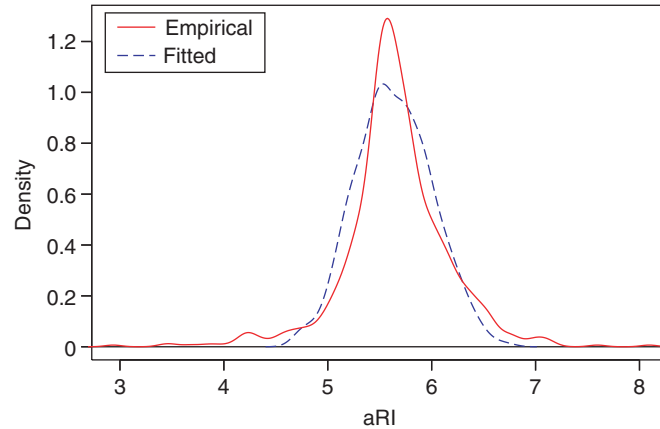
This represents, for all agents, the annual change in sentiment, and should be compared with the actual change in sentiment for each agent over that period. Positive τ would indicate an overall positive sentiment with respect to the M agents, and negative τ would indicate the opposite.

3.5.2 Network drag

The total influence of the system represents the effect on the sentiment of the network as a single entity. Each agent in the network experiences a network drag, given by (3.13), which gives the annualized total influence, covering a period of y years, as a proportion of the agent's cumulative excess sentiment. Using the notation of the previous section, we define the network drag $ND(j)$ for agent j as

$$ND(j) = \frac{100\tau}{CES(j)}. \quad (3.13)$$

This drag is up for an agent with poor reputation scores, and down for an agent with good reputation scores. It measures the extent to which an agent is affected by other agents.

FIGURE 2 Typical empirical sentiment (aRI) density with fitted normal distribution.

Normal fit. Mean = 5.651. Standard deviation = 0.359.

4 RESULTS

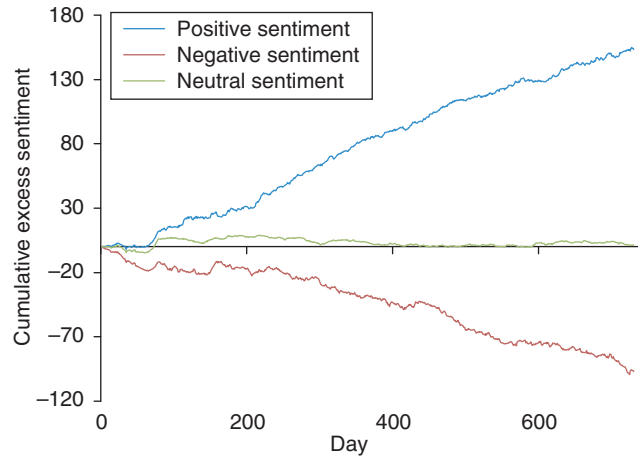
In this section, we give numerical results for the analyses in Section 3. Following some descriptive statistics for the reputation index, numerical values for the DeGroot influence matrix will be presented using the Bayesian methodologies described in Section 3. As a result, a network topology will be derived. In that network, it will be apparent that a robust consensus of sentiment is achieved. The consensus result will then be used to estimate a monetary value for reputation risk due to other agents, as opposed to the monetary value of an agent's own reputation.

4.1 Sentiment data description

Section 3.1.2 provided a brief description of the construction of the Alva Reputation Index. The data used in the calculations in this paper are the values of the aRI for ten UK banks (which are the "agents" of Section 3). They are not named for confidentiality reasons, and are therefore labeled Bank1, Bank2, ..., Bank10. For each of these banks, the aRI values have been assembled for the two-year period from January 2014 to December 2015 inclusive. The frequency is one aRI value per day for each bank, and no data is missing.

Figure 2 shows a typical sentiment distribution for one of the ten banks. A normal distribution can be fitted (using the maximum likelihood) to the empirical distribution, goodness-of-fit being determined by the transformed normal (TN-A) method described in Mitic (2015). The TN-A method was developed specially for fitting dis-

FIGURE 3 Typical cumulative excess sentiments showing positive, negative and neutral sentiment trends.

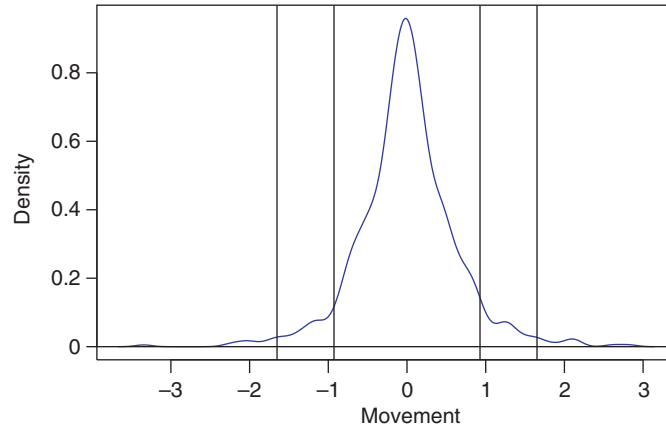


tributions encountered in operational risk, as it is independent of the size of the data set concerned. If the empirical mean and standard deviation of the sentiment data are m and s , respectively, then $N(m, (s/f)^2)$ provides a good normal fit, where f is a scale factor between 1 and 2. A good fit is obtained for $f \sim 1.4$. The precise distribution does not matter for the DeGroot analysis that follows. The mean sentiment is slightly in excess of the “par” value 5.5, indicating that sentiment is generally positive rather than negative.

Figure 3 shows typical plots of the daily cumulative excess sentiment (see Section 3.5.1) over a two-year period. These are derived from the profiles of the “best bank”, the “worst bank” and a “neutral bank” by perturbing each point on the profile by a small random amount. All resemble their originals closely (for reasons of confidentiality we cannot show the original data). The “positive” plot shows an almost linear increase in CES, and the “negative” plot shows a more volatile decrease. The “neutral” plot shows an almost constant daily sentiment of 5.5. An upward trending profile that then turns downward, or vice versa, has not been observed to date. From this (non)observation we conjecture that it is extremely difficult for a bank (or perhaps any other organization) to make good once it has obtained a poor reputation, and equally difficult for a bank’s reputation to deteriorate if it has been positive.

Figure 4 shows a typical histogram of movements, $M(i, t)$ (see (3.6)), of the aRI for one of the banks over the two-year period from January 2014 to December 2015.

FIGURE 4 Typical movements density, showing near-zero mean and small standard deviation.



Movements: mean = -0.00011 . Standard deviation = 0.633 .

The plot also shows the 5% and 1% tail cutoff points, which are discussed following the definition movement count in (3.7).

As expected, the profile in Figure 4 is similar to that shown in Figure 2. The differences $aRI_t - aRI_{t-1}$ are not independent: a correlation analysis shows that the aRI follows an autoregressive AR(1) statistical model (ie, the aRI value at any time t is significantly correlated with the aRI value at the previous time $t - 1$). The profile of $aRI_t - aRI_{t-1}$ fits a normal distribution (again using the maximum likelihood with the TN-A goodness-of-fit test in Mitic (2015)). Figure 4 shows an empirical standard deviation of 0.633 for the bank it represents, with a mean close to zero.

4.2 The influence matrix: numerical results


Setting precise numerical values for these limits is a matter of judgement, but defining L_H and L_{VH} (see (3.7)) in terms of percentages of all absolute movements provides a statistically objective way forward. Considering all movements, we designate 5% of the largest absolute movements as “large”, and the minimum “large” movement gives the threshold $L_H = 0.9276$. Similarly, we designate 1% of the largest absolute movements as “very large”, and the minimum “very large” movement gives the threshold $L_{VH} = 1.6516$. Setting higher values has the disadvantage that, as the limits increase, the number of items that correspond to those limits decreases to zero quite rapidly. Such “zero” results produce zero influence in most cases, as will be shown below. Setting lower values (eg, $L_H = 0.75$ and $L_{VH} = 0.95$) results in more nonzero

influences, and a fully connected influence network. Small variations in these limits have little effect on the calculated T . The sensitivity of T with respect to changes in L_H and L_{VH} is explored fully in Section 4.6.

The influence matrix T of (3.1), obtained using the data introduced in Section 4.1 and calculated using (3.8) and (3.9), is as follows:

$$T = \begin{pmatrix} 0.459 & 0.084 & 0 & 0 & 0 \\ 0.088 & 0.237 & 0 & 0 & 0 \\ 0 & 0 & 0.238 & 0 & 0.133 \\ 0 & 0 & 0.109 & 0.437 & 0 \\ 0 & 0 & 0.103 & 0.099 & 0.375 \\ 0.052 & 0.049 & 0 & 0 & 0 \\ 0.058 & 0.078 & 0 & 0.038 & 0 \\ 0.049 & 0.095 & 0.050 & 0 & 0.033 \\ 0 & 0 & 0 & 0.064 & 0.109 \\ 0.103 & 0.095 & 0 & 0 & 0 \\ & & & & & 0.079 & 0.115 & 0.189 & 0.074 & 0 \\ & & & & & 0.131 & 0.177 & 0.192 & 0.062 & 0.113 \\ & & & & & 0.122 & 0.133 & 0.179 & 0.194 & 0 \\ & & & & & 0.160 & 0.140 & 0.154 & 0 & 0 \\ & & & & & 0 & 0.146 & 0.181 & 0.096 & 0 \\ & & & & & 0.479 & 0.204 & 0.142 & 0.073 & 0 \\ & & & & & 0.086 & 0.518 & 0.084 & 0.097 & 0.041 \\ & & & & & 0.078 & 0.120 & 0.483 & 0.060 & 0.033 \\ & & & & & 0.162 & 0.084 & 0.117 & 0.465 & 0 \\ & & & & & 0 & 0 & 0.146 & 0.100 & 0.556 \end{pmatrix}. \quad (4.1)$$

An immediate observation from (4.1) is that there are 0 entries, implying that the graph associated with the weight matrix T is not fully connected. Not all agents are therefore able to influence all others directly. This is due to the stringent conditions on the values of L_H and L_{VH} . If these are relaxed, more influences result and a fully connected network is observed. Summing the values in each column of T reveals that the greatest column sums are in columns 6–9. This is an indication that these actors are the major influencers, an observation that will be confirmed when considering powers of T . Figure 5 shows the network represented by T . The directional edges from agent X to agent Y should be taken to mean “ X is influenced by Y ”. The thickness of an edge is an indication of the magnitude of the influence (thin = noninfluential, thick = influential).

FIGURE 5 Influence network showing degrees of influence and dominant agents 6–9.

Thick edges indicate high influence; thin edges indicate low influence.


Most edges in Figure 5 are bidirectional: pairs of agents are influenced by each other. However, the influence between any given pair is not necessarily symmetric. This is apparent from (3.10): T is nonsymmetric. The accumulation of edges terminating at agents 6, 7, 8 and 9 indicates that these are the dominant agents, as will be shown in the following section.

4.3 Powers of the influence matrix and consensus: numerical results

The matrix T^2 represents all influences resulting from paths of length 2 in the network represented by T . All entries in T^2 are nonzero, and the network corresponding to T^2 is fully connected. Therefore, all agents are able to influence all other agents either directly or indirectly via a path of length 2. The result from Berger (1981), mentioned in Section 3.3.1, then guarantees convergence of powers of T , and consensus. In this case the result is very strong. Only one column of any power of T is required to fulfill the conditions of Berger's theorem. For the lowest power of T possible, all columns contain strictly positive elements.

Higher powers of T begin to reveal increasing consistency in the column values. In the matrix T^6 , the values in any given column are identical to two significant figures. The column values in T^{10} are identical to three significant figures, and at this stage convergence to a consensus is apparent. This result echoes the finding of Gai and Kapadia (2010), who also observed very fast convergence. They detected convergence having calculated the fourth power of their influence matrix, indicating

FIGURE 6 Influence network for T_∞ , rounded down to the nearest first decimal place, showing partition of influencers (6,7,8,9) and influenced (1,2,3,4,5,10).



Thick edges indicate high influence; thin edges indicate low influence.

that contagion had spread fully. Equation (4.2) shows the first and last rows in T^{12} . All other rows in T^{12} (ten in total, including the two shown) are identical, to three significant figures. The rate of convergence depends on the number of agents. Matrix T^{12} can be taken as representative of $T_\infty = \lim_{t \rightarrow \infty} (T^t)$, with the elements shown being correct to three significant figures:

$$T_\infty = \begin{pmatrix} 0.078 & \dots & 0.078 \\ 0.072 & \dots & 0.072 \\ 0.024 & \dots & 0.024 \\ 0.036 & \dots & 0.036 \\ 0.038 & \dots & 0.038 \\ 0.152 & \dots & 0.152 \\ 0.212 & \dots & 0.212 \\ 0.207 & \dots & 0.207 \\ 0.127 & \dots & 0.127 \\ 0.053 & \dots & 0.053 \end{pmatrix}. \tag{4.2}$$

It is readily apparent that the elements in rows 6–9 in T_∞ are much larger than the elements in the others. Agents 6–9 are therefore dominant influencers. Further, agents 7 and 8 are approximately twice as influential as agents 6 and 9. Agents 3–5 are the least influential. Numeric changes to parameters in DeGroot methodology result in the same conclusions, indicating that the assumptions made are robust.

The network represented by T_∞ in (4.2) is fully connected. To make better sense of it, it is instructive to round down all the figures in (4.2) to the first decimal place. Each row in T_∞ is then

$$(0, 0, 0, 0, 0, 0.1, 0.2, 0.2, 0.1, 0). \quad (4.2')$$

The entries in (4.2') emphasize the dominance of agents 6–9. A transition matrix comprising ten rows identical to the single row in (4.2') corresponds to the influence network in Figure 6. There is a direct path to agents 6, 7, 8 and 9 from all of the others, and also from agents 6, 7, 8 and 9 to themselves. The network is partitioned into two blocks. The influencers are agents 6, 7, 8 and 9 and the influenced are agents 1, 2, 3, 4, 5 and 10. The thickness of the arcs in Figure 6 indicates the degree of influence (thick, 0.2 indicates “most influence”; thin, 0.1 indicates “not as influential”).

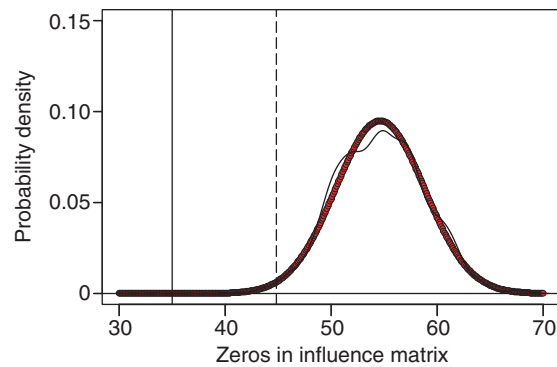
The results in (4.2) are surprising. Agents 6–9 are not organizations that are constantly the subject of headlines in the press. Rather, they quietly score well in league tables (eg, “best buy” tables), and have a general reputation (in a loose sense!) for quality of service. The least influential, agent 3, has consistently enjoyed a reputation as “best bank” for more years than the two in this analysis. By way of explanation, we can make some conjectures, but we cannot substantiate them without further analysis.

- The way to gain a good reputation is to steer clear of negative headlines. This comment is motivated by the upward CES sloping profile of agent 3, which resembles the upward sloping profile in Figure 3. This bank is not an influencer.
- Agents 2, 5 and 10 are banks that have suffered negative sentiment as a result of, eg, illegal dealings, regulator fines, technical outages and deficient conduct. They are (relatively) benign influencers. A possible explanation is that the negative sentiment for agents 2, 5 and 10 is due to matters that do not directly affect retail customers, even though that negative sentiment might have a very serious origin. The CES profiles for agents 2 and 5 resemble the downward sloping profile in Figure 3.
- The cumulative excess sentiment profile for agent 10 is a curiosity. It resembles the flat profile in Figure 3. Although this bank was subject to huge regulatory fines during the two years of this study, it has also been innovative technologically. A possible explanation is that negative sentiment and positive sentiment have canceled each other out. Bank 10 is not an influencer.

4.3.1 *Co-movements of individual reputations: randomized trial tests*

The results presented in preceding sections are calculated using the unique data set available, leading to the unique influence matrix T (see (3.13)). The significance

FIGURE 7 Comparison of number of zeros for randomly generated co-movements (black curve: randomized co-movements; red curve: fitted normal distribution) with number of zeros in the influence matrix (solid vertical line) and 1% significance value (dashed vertical line).



of T may be assessed by randomizing the co-movements used to formulate T , and defining an appropriate metric with which to do the assessment. The co-movements are tuples of movements $\{M(i, t_1), M(i, t_2)\}$ for agent i , defined in (3.6), where t_1 and t_2 are different days. Randomizing co-movements may be used to test if the actual set of movements is in any way different from any other set of movements. A convenient metric to apply is to count the number of zero entries in T , which is 35. Recalculating the post-randomization influence matrix several times quickly establishes that the number of zero entries in any post-randomization influence matrix is nearer to 55. Further, generating n (> 100) post-randomization influence matrixes $\{T_{P_1}, T_{P_2}, \dots, T_{P_n}\}$ indicates that $\nu(T_P)$, the number of zero entries in an influence matrix (either actual or post-randomization) T_P , has an approximately normal distribution. These observations appear to reveal that the T of (4.1) is indeed unlike any influence matrix resulting from random co-movements.

A formal approach, in which 1000 random co-movements were generated, yielded the following results for $\nu(T_P)$:

n	1000
Fitted normal distribution	$\nu(T_P) \sim N(54.645, 4.218^2)$
Median $\nu(T_P)$	55
Maximum $\nu(T_P)$	67
Minimum $\nu(T_P)$	40

Figure 7 shows a density plot of the random co-movements and a density plot of the fitted normal distribution. It demonstrates clearly that the number of zeros for the actual influence matrix (4.1) is far into the left tail of the distribution. In terms of a formal significance test, testing the null hypothesis $\nu(\mathbf{T}_P) = 35$ against the alternative hypothesis $\nu(\mathbf{T}_P) > 35$ using a z -test based on the fitted normal distribution shows that the measured mean, 54.645, differs very significantly from the actual mean, 35 (p -value 4×10^{-6} , 4.42 “standard deviations”). Consequently, the measured influence matrix can be taken to represent a nonrandom effect of reputation contagion.

4.3.2 The effect of autocorrelations on the influence matrix

The values of the aRI are potentially autocorrelated, so special treatment may be needed in the calculation of the influence matrix entries in (3.8) and (3.9). A positive autocorrelation would be expected to reduce the values of the movements $M(i, t)$ of (3.6), with a commensurate reduction in the values of L_H and L_{VH} to give reasonable values for \mathbf{T} (see the discussion immediately after (3.7) and the results reported in Section 4.2). The entries in \mathbf{T} depend on a ratio of counts, which is expected to be largely unchanged by a reduction in movements. Therefore, autocorrelations were not used to derive the result of Section 4.2.

To confirm that autocorrelations have a statistically non-significant effect on the structure of the consensus network, autocorrelation coefficients r_i ($i = 1, \dots, 10$) were calculated for all ten banks. They were all positive, with a mean of 0.34. Recalculating \mathbf{T} and \mathbf{T}_∞ using $(1 - r_i) \times M(i, t)$, instead of $M(i, t)$ (in (3.6)), and scaling the 1% and 5% thresholds in Section 4.2 by a factor $1 - 0.34 = 0.66$, produced a reconstructed consensus network with the same properties as the network in Figure 6.

4.4 Consensus on opinion: subjective initial opinion vector

In this section, we discuss the formulation of the initial opinion vector \mathbf{p} , introduced in Section 3.3. Equation (3.3) gives the consensus opinion, $\mathbf{p}(\infty)$, in terms of an initial opinion, $\mathbf{p}(0)$. The values of the entries in $\mathbf{p}(0)$ could be formulated in any acceptable way, and in this section we give examples in which $\mathbf{p}(0)$ is formulated objectively, and also by a loose concept of “reputation”. It should be stressed that values assigned to $\mathbf{p}(0)$ could be purely subjective, or they could be based on evidence from, eg, opinion polls or surveys. The examples given here originate from subjective considerations, although those considerations are backed by their cumulative sentiment profiles, such as those shown in Figure 3.

The previous section referred to three banks (Bank2, Bank5 and Bank10) that have had very negative sentiment, and one that has had the opposite (Bank3). These two categories of bank have, respectively, cumulative sentiment profiles very similar to the “negative” and “positive” profiles in Figure 3. These profiles demonstrate the loose

concept of “reputation”, and can be defined by formulating the appropriate vectors $\mathbf{p}(0)$. Indeed, Pan (2010) uses arbitrary values in $\mathbf{p}(0)$, and Pan (2012) uses normally distributed values in $\mathbf{p}(0)$. Neither is appropriate in the context of opinions on the banking sector. In the current context, the result of $\mathbf{T}_\infty(\mathbf{p}(0))$ clearly depends on $\mathbf{p}(0)$, since \mathbf{T}_∞ tends to a limit as r tends to infinity. So it seems hard to justify using a random $\mathbf{p}(0)$, or any other $\mathbf{p}(0)$ that is not directly related to an observable.

In the case of three negative perceptions of sentiment, the normalized vector $\mathbf{p}(0)$ is given by $\mathbf{p}(0) = (0, \frac{1}{3}, 0, 0, \frac{1}{3}, 0, 0, 0, 0, \frac{1}{3})$. This says that agents 2, 5 and 10 are equally perceived as having a particular reputation: poor in this case. From (3.3), the consensus view, $\mathbf{p}(\infty)$ is a vector in which all entries are 0.0547, correct to three significant figures. The consensus view is therefore that all agents are equally “bad”. Agents 2, 5 and 10 are not as “bad” as was initially proposed, and the others are not perfectly “good”.

For the case of agent 3, who has what is perceived to be a very positive reputation, we set $\mathbf{p}(0) = (0, 0, 1, 0, 0, 0, 0, 0, 0, 0)$. The consensus view in this case, $\mathbf{p}(\infty)$, is a vector in which all entries are 0.0238, correct to three significant figures. The interpretation of this consensus view is that there is some “good” in all of the agents, and the positive sentiment attributed to agent 3 is much reduced.

4.4.1 Consensus of opinion: objective initial opinion vector

The results in this section are for the theory in Section 3.5.1. The intention is to provide a robust assumption of initial perception.

To provide a more objective view on perceptions of “good” and “bad” banks, the CES measure is a very effective basis because it is visually striking. Equation (3.11) gives the CES measure \mathbf{F} , for a set of agents. Calculating \mathbf{F} for the ten banks in this study over the two years 2014–15, with the “par” value $\overline{\text{aRI}}_j = 5.5$ for $j = 1, \dots, 10$, shows that all elements in \mathbf{F} are in the range $(-300, 300)$. Specifically, they are

$$\mathbf{F} = (111.3, -212.3, 195.7, 47.4, -198.7, 133.2, -64.0, 76.3, 156.1, -27.9). \quad (4.3a)$$

With a linear map of the range $(-300, 300)$ to $(0, 1)$, the entries in \mathbf{F} produce, after normalizing, the initial perception vector (4.3a). Correct to three significant figures, this gives

$$\mathbf{p}(0) = (0.128, 0.027, 0.154, 0.180, 0.031, 0.135, 0.073, 0.117, 0.141, 0.085). \quad (4.3b)$$

All the resulting entries in the corresponding $\mathbf{p}(\infty)$ (using (3.3)) are 0.104. This value is then the one used in all ten entries of vector \mathbf{J} of (3.12), so that $\mathbf{J} = \{0.104, 0.104, \dots, 0.104\}$.

TABLE 1 Network drag percentages.

Bank j	ND(j)
1	10.1
2	-5.3
3	5.8
4	23.8
5	-5.7
6	8.5
7	-17.6
8	14.8
9	7.2
10	-40.5

4.4.2 Consensus of opinion: total influence

With the preceding value for \mathbf{J} and $y = 2$ years, the value of the total influence defined in (3.12) is then

$$\tau = \frac{\mathbf{\Gamma} \cdot \mathbf{J}}{y} \sim 11.3. \quad (4.4)$$

This consensus value is an effective “smoothing” of the initial perception vector, and should be seen in the context of the range of values in vector $\mathbf{\Gamma}$ (see (4.3a)). It represents a consensus view that sentiment expressed toward the ten banks is positive rather than negative, but only just. Perceptions for agents initially perceived as “bad” are boosted, and agents initially perceived as “good” are muted. The value 11.3 is a mediocre total influence, and the conclusion that, overall, banks enjoy a positive reputation may be surprising.

4.4.3 Consensus of opinion: network drag

With the preceding value for τ , the CES measures $\mathbf{\Gamma}$ (from (4.3a)), the values of the network drag defined in (3.13) for all ten banks are given in Table 1.

A bank that has a cumulative excess sentiment that is nearer zero can attribute all or most of its cumulative sentiment to network effects. For any bank, the actual cumulative score is the sum of an intrinsic score, earned by the bank itself, and a score due to network contagion.

4.5 Stubborn agents

A stubborn agent is one who values its own opinion much more highly than the opinion of others. The basis of the procedure for modeling stubborn agents, the Friedkin amendment to the DeGroot model, was briefly discussed in Section 3.3.2.

Using (3.5), the limiting matrix $(\mathbf{I} - \mathbf{A}\mathbf{T})^{-1}(\mathbf{I} - \mathbf{A})$ exists but, as might be expected, does not have the same values in each column. Indeed, a sequence defined by (3.4) would not necessarily converge. Using the initial perception vector (4.3b), the limiting values $\mathbf{p}(\infty)$ are

$$\mathbf{p}(\infty) = (0.118, 0.042, 0.144, 0.110, 0.053, 0.122, 0.084, 0.110, 0.128, 0.091). \quad (4.5)$$

The rank order of the elements in (4.5) is the same as the rank order of the elements in the initial perception vector (4.3b), although the ratio of pairs of elements is different. The rank order indicates the influence that an agent has on other agents, so if the rank orders of agents in the “stubborn” and the “non-stubborn” case are the same, agents are already quite stubborn. This indicates that agents stick to their own views overall, but are very marginally influenced by others.

4.6 Sensitivity to large sentiment movements

The analysis in Section 4.2, with the values in matrix (3.13), the network in Figure 5 and all numerical results that follow, was derived using particular thresholds L_H and L_{VH} . In this section, we vary the values of these parameters and investigate the results. It makes sense to concentrate on changes in daily sentiment value that are, in the loose sense of the word, “large”. Smaller changes are essentially noise. Therefore, L_H and L_{VH} are both varied between 0.8 and 2.2. Approximately 21.4% of absolute changes exceed 0.8, and approximately 0.70% of absolute changes exceed 2.2. Using smaller or larger values provides no further insight.

In order to compare results, we use the mean of nonzero entries in the influence matrix, $\mu(\mathbf{T})$, to give a single-figure measure of each influence matrix that results from each pair of thresholds (L_H, L_{VH}) . Specifically,

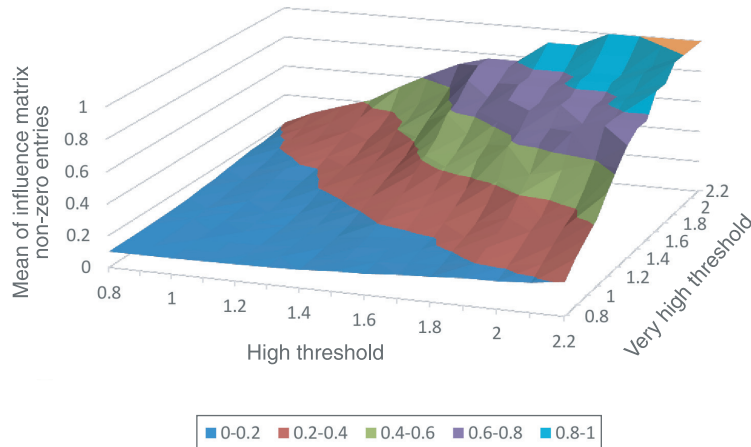
$$\mu(\mathbf{T}) = \frac{\sum_{i,j=1}^n (\mathbf{T}_{ij} \mid \mathbf{T}_{ij} > 0)}{\#(\mathbf{T}_{ij} \mid \mathbf{T}_{ij} > 0)}, \quad (4.6)$$

where $n = 10$, and function $\#(\cdot)$ denotes the number of elements in the argument.

This measure has the advantage that it gives an indirect indication of the number of nonzero entries in the influence matrix and their strength (recall that the row sums in the nonzero entries are all 1, so that the mean of all entries in a 10×10 influence matrix is $10 \times 1/100 = 0.1$). Figure 8 shows a surface plot of $\mu(\mathbf{T})$ against the two thresholds L_H and L_{VH} .

The general trend in Figure 8 is clear: $\mu(\mathbf{T})$ increases with increasing L_H and L_{VH} . At the “peak” ($L_H \sim 2.2$ and $L_{VH} \sim 2.2$), a small plateau is discernible. In this region, the influence matrix is, or closely resembles, the 10×10 identity matrix. All agents influence only themselves. The conditions to measure influence are too stringent to

FIGURE 8 The effect of variation of Bayesian thresholds showing L_H and L_{VH} on the horizontal plane and the metric $\mu(T)$ vertically.



Greatest “high” and “very high” thresholds: agents influence only themselves. Least “high” and “very high” thresholds: agents influence all others. Mid-value “high” and “very high” thresholds: agents influence some others.

be of use: the conditional probabilities in (3.8) are zero. When L_H and L_{VH} are both small, the opposite applies. The conditions to measure influence are sufficiently weak to allow all agents to influence all others. The influence matrix has no zero entries and the corresponding influence network is fully populated. The “small L_H and L_{VH} ” case applies approximately for the domain $\{(H, VH) : H > 0; VH > 0; H + VH < 2.2\}$, which is the region in Figure 8 that looks largely flat. Outside that region, values for L_H and L_{VH} may be set such that they provide a reasonable indication of a rare event linked to another rare event. The values given in Section 4.2 do this. The resulting influence matrix is robust with respect to changes in the values of L_H and L_{VH} , provided that they are not “small L_H and L_{VH} ” cases. Powers of the influence matrix converge, usually quickly, to the same limiting case, matrix (4.1).

5 MONETARY GAIN/LOSS DUE TO REPUTATION CONTAGION

In this section, we report results that relate reputation contagion effects, propagated through the network described in Section 4, to actual monetary losses or gains, expressed as sales.

In Mitic (2017), we reported the annual effect of sentiment on annual sales. In particular, the effects of sustained positive and sustained negative sentiment were

TABLE 2 Expected values of the effect of sentiment on product sales.

Product	Positive sentiment (%)	Negative sentiment (%)
Expected sales volume	1.6	2.3
Expected profit after tax	0.7	0.9
Stressed sales volume	3.4	7.9
Stressed profit after tax	1.3	3.6

discussed. The case of the Volkswagen diesel emissions falsification (Society of Motor Manufacturers and Traders 2015) is an example of sustained negative sentiment. In that case, a change in the aRI from 6.20 to 2.10 over a one-week period could be linked to a year-on-year drop of 19.99% in sales of new cars (3238 vehicles). Cases of sustained positive sentiment are hard to identify, but some indications are provided in Winer (2001).

The values reported in Mitic (2017) for the changes in product sales, income and profit, given as a percentage of total sales, income and profit, are summarized in Table 2. The expected changes were calculated using fitted probability distributions, and the stressed changes were calculated using additional scenarios.

The result in Table 1 for Bank1 (approximately 10% network drag) is applicable to all of the values in Table 2. The interpretation of the effect of network drag in this case is that the negative sentiment figures are not as high as those quoted in Table 2, and the positive sentiment figures are higher than those quoted in Table 2: the network has dragged them both down.

6 DISCUSSION

6.1 General conclusions

The aim of this study was to assess the effects of the reputation of the members of a group on any single member of the group using the concepts of social influence and convergence in belief. The principal effect was to find what proportion of an agent's reputation is attributable to other agents in a network. The DeGroot model was used to calculate a consensus view of reputation. One way to think about how this comes about is to note the precise meaning of the term "agent", which has been used extensively in this paper. It is tempting to think of "agent" as referring to a singular entity, such as a bank. Actually, it refers to a collection of that bank's customers, any one of which may potentially be influenced by any other. An individual customer may, indeed, be the customer of several banks, so the bank-agent association is unlikely to be disjoint. It is assumed that the DeGroot analysis is not affected. After an exchange

of views they reach a consensus, which is represented by powers of the influence matrix.

Influence matrixes do not always converge when they are repeatedly multiplied by themselves. In this study, the influence matrix does converge, and this convergence is apparent after only six iterations. The interpretation of this result is that a chain of customers, each talking to the next customer in the chain about their bank, need only be six customers long before nearly every customer has had their say. The degree of influence is determined by setting threshold parameters so that a large influence on the customers of one bank has an even larger influence on the customers of another bank.

Previous results indicated that it was possible to link sentiment with monetary values expressed in terms of sales. Given a value for the change in sales that can be attributed to a change in sentiment, it is possible to allocate a proportion of those sales to network sentiment. The precise proportion depends on the absolute value of the cumulative excess sentiment and the total influence of agents as a group.

The principal numerical finding for this study is the extent of the “network drag” in Table 1. Banks that have an upward trending cumulative excess sentiment profile (similar to the positive sentiment profile in Figure 3) are subject to a downward “network drag” of 5–10%. Banks 1, 3, 6 and 9 are of this type. They are above-average performers in terms of reputation, and their cumulative excess sentiment profiles are depressed due to the presence of other banks. Conversely, banks 2 and 5 experience an upward “network drag” of about 5%. Their reputations are poor, and are inflated due to the presence of other banks. In the middle are banks that have high “network drags”, either up or down (banks 4 and 10 in particular). They are average performers in terms of reputation, and the high “network drag” values indicate that a large percentage of their reputation is attributable to other banks. The most influential banks (banks 6–9; see Figure 6) have intermediate valued “network drags”. The total influence value, 11.3 (see (4.4)), provides an overall measure of sentiment for the network. It is positive, indicating that the agents in the network have an overall positive reputation. This observation is reinforced by the sum of “network drag” values, 1.1 (from Table 1). Starting with the objective initial opinion vector (see (4.3b)), the calculated value of $\mathbf{p}(\infty)$ is 0.104 (see the end of Section 4.4.1). This result further amplifies the “network drag” effect: opinions converge to a consensus that, in this case, is close to the mean initial opinion.

6.2 Further work

Three principal items of further work follow immediately from the current analysis. The first is to investigate the time dependence of the influence matrix, and this work is already well advanced. Preliminary results indicate that such influence matrixes show

a reasonable degree of consistency, but outliers do occur. The second is to introduce a shock into the system in order to examine the effect on the network of extreme events for any one actor. The third is to consider an alternative sentiment-propagation model. In particular, the voter model, in which agents make decisions based on the views of their immediate neighbors, might yield interesting results.

DECLARATION OF INTEREST

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

ACKNOWLEDGEMENTS

I am grateful for the support of Alva Group for granting permission to use their data and discuss their methodology, and for their interest and assistance in the preparation of this paper.

REFERENCES

- Aldasoro, I., delli Gatti, D., and Faia, E. (2015). Bank networks: contagion, systemic risk and prudential policy. Working Paper 5182, SAFE (<http://doi.org/bv98>).
- Battiston, S., di Iasio, G., Infante, L., and Pierobon, F. (2013). Capital and contagion in financial networks. IFC Bulletin 39, Bank for International Settlements (<http://bit.ly/2h9j26Z>).
- Beale, N., Rand, D. G., Battey, H., Croxson, K., May, R. M., and Nowak, M. A. (2011). Individual versus systemic risk and the regulator's dilemma. *Proceedings of the National Academy of Sciences* **108**(31), 12 647–12 652 (<http://doi.org/c83nss>).
- Berger, R. L. (1981). A necessary and sufficient condition for reaching a consensus using DeGroot's method. *Journal of the American Statistical Association* **76**, 415–418 (<http://doi.org/bwbb>).
- Brown, M., Trautmann, S. T., and Vlahu, R. (2014). Understanding bank-run contagion. Working Paper 1711, European Central Bank (<http://bit.ly/2h11dDS>).
- Bruder, M., Fischer, A., and Manstead, A. (2014). Social appraisal as a cause of collective emotions. In *Collective Emotions*, Scheve, C., and Salmella, M. (eds), Chapter 10. Oxford University Press.
- Chan-Lau, J. A., Mitra, S., and On, L. L. (2007). Contagion risk in the international banking system and implications for London as a global financial center. Working Paper WP/07/74, International Monetary Fund, Washington, DC (<http://doi.org/bwbc>).
- Chatterjee, S., and Seneta, E. (1977). Towards consensus: some convergence theorems on repeated averaging. *Journal of Applied Probability* **14**(1), 89–97 (<http://doi.org/bwbd>).
- Chaudhari, M., and Govilkar, S. (2015). A survey of machine learning techniques for sentiment classification. *International Journal on Computational Science and Applications* **5**(3), 13–23 (<http://bit.ly/2h12NFS>).

- Collomb, A., Costea, C., Joyeux, D., Hasan, O., and Brunie, L. (2014). A study and comparison of sentiment analysis methods for reputation evaluation. Technical Report RR-LIRIS-2014-002, Laboratoire d'Informatique en Image et Systèmes d'Information Lyon (<http://bit.ly/2hbAHwv>).
- Das, A., Gollapudi, S., and Munagala, K. (2014). Modeling opinion dynamics in social networks. In *Proceedings of the 7th ACM International Conference on Web Search and Data Mining*. ACM Publications, New York (<http://doi.org/bwbq>).
- Davies, M. (2010). The corpus of contemporary American English as the first reliable monitor corpus of English. *Literary and Linguistic Computing* **25**(4), 447–65 (<http://doi.org/d7b5ms>).
- DeGroot, M. H. (1974). Reaching a consensus. *Journal of the American Statistical Association* **69**, 118–121 (<http://doi.org/bwbs>).
- Degryse, H., and Nguye, G. (2007). Interbank exposures: an empirical examination of contagion risk in the Belgian banking system. *International Journal of Central Banking* **3**(16), 123–171 (<http://bit.ly/2hUump>).
- DeMarzo, P., Vayanos, D., and Zwiebel, J. (2003). Persuasion bias, social influence and unidimensional opinions. *Quarterly Journal of Economics* **118**, 909–968 (<http://doi.org/fnw3jv>).
- Dungeya, M., and Gajurela, D. (2015). Contagion and banking crisis: international evidence for 2007–2009. *Journal of Banking and Finance* **60**, 271–283 (<http://doi.org/bwbt>).
- Egloff, E. (2015). Affaire Volkswagen: le risque de contagion inquiète le monde de l'automobile. *Le Figaro*, September 24 (<http://bit.ly/2i9F775>).
- Elliott, M., Golub, B., and Jackson, M. O. (2014). Financial networks and contagion. *American Economic Review* **104**(10), 3115–3153 (<http://doi.org/bwbv>).
- Federal Reserve Bank of Boston (1995). Rating the adequacy of risk management processes and internal controls at state member banks and bank holding companies. Supervisory Letter SR 95-51 (SUP), Boston Fed (<http://bit.ly/2hqV1T>).
- Fiordelisi, F., Soana, M.-G., and Schwizer, P. (2013). The determinants of reputational risk in the banking sector. *Journal of Banking and Finance* **37**(5), 1359–1371 (<http://doi.org/bwbx>).
- Fiordelisi, F., Schwizer, P., and Soana, M.-G. (2014). Reputational losses and operational risk in banking. *European Journal of Finance* **20**(2), 105–124 (<http://doi.org/bwbw>).
- Fombrun, C. J., Ponzi, L., and Newburry, W. (2015). Stakeholder tracking and analysis: the RepTrak system for measuring corporate reputation. *Corporate Reputation Review* **18**, 3–24 (<http://doi.org/bwbz>).
- Forney, G. D. (1973). The Viterbi algorithm. *Proceedings IEEE* **61**(3), 268–278 (<http://doi.org/b7hr4m>).
- Friedkin, N. E. (2011). A formal theory of reflected appraisals in the evolution of power. *Administrative Science Quarterly* **56**(4), 501–529 (<http://doi.org/bwb3>).
- Friedkin, N. E., and Johnsen, E. C. (1990). Social influence and opinions. *Journal of Mathematical Sociology* **15**(3), 193–205 (<http://doi.org/bghfwr>).
- Frobenius, G. (1912). *Über Matrizen aus nicht negativen Elementen*, pp. 456–477. Königliche Gesellschaft der Wissenschaften, Gottingen (<http://doi.org/bwb2>).
- Gai, P., and Kapadia, S. (2010). Capital and contagion in financial networks. Working Paper 383, Bank of England (<http://bit.ly/2hqFI03>).

- Georg, C.-P. (2011). The effect of the interbank network structure on contagion and common shocks. Discussion Paper 12/2011, Deutsche Bundesbank (<http://bit.ly/2hqzXzF>).
- Georg, C.-P. (2014). Contagious herding and endogenous network formation in financial networks. Working Paper 1700, July, European Central Bank (<http://bit.ly/2hU7hIQ>).
- Ghaderi, J., and Srikant, R. (2012). Opinion dynamics in social networks: a local interaction game with stubborn agents. In *Proceedings of the American Control Conference*. IEEE Press, New York (<http://doi.org/bwb4>).
- Golub, B., and Jackson, M. O. (2010). Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics* **2**(1), 112–149 (<http://doi.org/dfpqg4>).
- Grabisch, M., and Rusinowska, A. (2011). A model of influence based on aggregation functions. *Mathematical Social Sciences* **66**(3), 316–330 (<http://doi.org/bwb5>).
- Hałaj, G., and Kok, C. (2014). Modeling emergence of the interbank networks. Working Paper 1646, European Central Bank (<http://bit.ly/2hGXgll>).
- Herrera-Viedma, E., Herrera, F., and Chiclana, F. (2002). A consensus model for multi-person decision making with different preference structures. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans* **32**(3), 394–402 (<http://doi.org/b7qvqt>).
- Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.
- Jackson, M. O., and Lopez-Pintado, D. (2013). Diffusion and contagion in networks with heterogeneous agents and homophily. *Network Science* **1**(1), 49–67 (<http://doi.org/bwb6>).
- Jia, P., MirTabatabaei, A., Friedkin, N. E., and Bullo, F. (2015). Opinion dynamics and the evolution of social power in influence networks. *SIAM Review* **57**(3), 367–397 (<http://doi.org/bwb7>).
- Jurafsky, D., and Martin, J. H. (2008). *Speech and Language Processing*. Prentice Hall, Englewood Cliffs, NJ (<http://stanford.io/2hH5EaS>).
- Kučera, H., and Francis, W. N. (1967). *Computational Analysis of Present-Day American English*. Brown University Press, Providence, RI.
- Liu, B. (2015). *Sentiment Analysis: Mining Opinions, Sentiments and Emotions*. Cambridge University Press (<http://doi.org/bwb8>).
- Londono, J. M. (2016). Bad bad contagion. Discussion Paper 1178, Federal Reserve Bank International Finance (<http://doi.org/bwb9>).
- Marcus, M., Antorini, B. S., and Marcinkiewicz, M. A. (1993). Building a large annotated corpus of English: the Penn Treebank. Technical Report MS-CIS-93-87, University of Pennsylvania Department of Computer and Information Science (<http://bit.ly/2i03gJZ>).
- Mitic, P. (2015). Improved goodness-of-fit measures. *The Journal of Operational Risk* **10**(1), 77–126 (<http://doi.org/bwcb>).
- Mitic, P. (2017). Reputation risk: quantified. In *Complex Systems 2017: Proceedings of the 3rd Complex Systems Conference*, Rzevski, G., and Brebbia, C. (eds). Wessex Institute Transactions. WIT Press, Southampton.
- Newman, M. (2010). *Networks: An Introduction*. Oxford University Press (<http://doi.org/cfnzxsx>).
- Organisation for Economic Co-operation and Development (2012). Financial contagion in the era of globalised banking. Policy Notes 14, Economics Department, OECD (<http://bit.ly/2h12pai>).

- Owram, K. (2015). The Volkswagen contagion: how the emissions-rigging scandal is spreading to Canadian firms. *Business Financial Post*, September 24 (<http://bit.ly/2ieqTxE>).
- Pan, Z. (2010). Trust, influence, and convergence of behavior in social networks. *Mathematical Social Sciences* **60**, 69–78 (<http://doi.org/br7tkr>).
- Pan, Z. (2012). Opinions and networks: how do they effect each other. *Computational Economics* **39**(2), 157–171 (<http://doi.org/fd8nhf>).
- Pan, Z., and Gilles, R. P. (2010). Naive learning and game play in a dual social network framework. Economics Working Paper 10-01, Queen's Management School, Queen's University Belfast (<http://bit.ly/2hUnUhy>).
- Pang, B., Lee, L., and Vaithyanathan, S. (2002). Thumbs up? Sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing*, Haji, J., and Matsumoto, Y. (eds), Volume 10, pp. 79–86. Association for Computational Linguistics, Stroudsburg, PA (<http://doi.org/c7z4vt>).
- Peleja, P., Santos, J., and Magalhães, J. (2014). Reputation analysis with a ranked sentiment-lexicon. In *Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1207–1210. ACM Publications, New York (<http://doi.org/bwcf>).
- Perron, O. (1907). Zur Theorie der Matrizes. *Mathematische Annalen* **64**(2), 248–263 (<http://doi.org/bzkkzw>).
- Perry, J., and de Fontnouvelle, P. (2005). Measuring reputational risk: the market reaction to operational loss announcements. Report, Federal Reserve Bank of Boston (<http://bit.ly/2i06jC1>).
- Ponzi, L. J., Fombrun, C. J., and Gardberg, N. A. (2011). RepTrak Pulse: conceptualizing and validating a short-form measure of corporate reputation. *Corporate Reputation Review* **14**(1), 15–35 (<http://doi.org/fpvnr>).
- Society of Motor Manufacturers and Traders (2015). New car registrations return to growth in November. Article, December 4, SMMT (<http://bit.ly/2hqDeie>).
- Tabak, B. M., and de Castro Miranda, R. C. (2013). Contagion risk within firm–bank bivariate networks. Working Paper 322, Banco Central do Brasil (<http://bit.ly/2hZOMtF>).
- Viterbi, A. J. (1967). Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory* **13**(2), 260–269 (<http://doi.org/dsmdmk>).
- Winer, R. S. (2001). A framework for customer relationship management. *California Management Review* **43**(4), 89–105 (<http://doi.org/bwccg>).