

Risk analysis of supply: comparative performance and short-term prediction

L. Walls, J. Quigley, M. Parsa & E. Comrie

Department of Management Science, University of Strathclyde, Glasgow, Scotland

ABSTRACT: Managing risk in the supply of inbound parts for the assembly of complex engineering systems is a key activity for manufacturing companies to avoid production delays and disruption. We present novel analysis of empirical data typically available in enterprise resource planning systems. The information generated through such analysis can provide useful information to support decisions to better manage supply risk. For example, identifying suppliers at risk of poor performance under different loadings or relative to peers, and assessing the financial value of investing in learning and developing suppliers. Our analysis is based on new methods developed for robust ranking, valuation of due diligence as well as more established regression modelling principles. We focus upon the ways our methods can be used to support supply risk decisions by showing how they have been applied in an industry case and the implications of our findings to date.

1 INTRODUCTION

1.1 *Motivating industry problem*

Our study has been driven by challenges facing supply managers in several, similar manufacturing companies. The managers have responsibilities that span the life of a supplier, from procuring, developing, maintaining and/or exiting the relationship. The manufacturing companies are all responsible for the assembly of engineering systems for which the sub-assemblies and component parts are sourced from suppliers globally. The design authority for parts can reside with the manufacturer, the supplier or both.

The set of suppliers for the manufacturing companies are continually in a state of change, with suppliers being exited and new ones being added. In addition, the volume of work apportioned to each supplier changes over time depending on projects commissioned by end customers for bespoke engineering systems. The manufacturers produce highly complex products which are customised to varying degrees. The customisation together with limited holding space and resource implications means it is not possible to stock a spare of every type of part. Therefore, parts are ordered as needed.

On-time delivery of parts from suppliers to the required functional and quality standard is necessary for the production of the final assembly so that the manufacturer meets end customer deadlines. However problems such as production delays and disruption arise when parts are delivered late or do not conform to requirement.

Heuristically the supply managers have an understanding of the causes and modes of supplier failure to deliver parts to the correct specification at the right time. However, the managers do not have models to support analysis that allows them to articulate and evidence the relationship between factors that drive supply risk and hence no reliable means of making prediction, even for relatively short-term planning horizons. The consequence is that there tends to be a reactive response to events that delay or disrupt supply. Our managers would like analytical models to provide them with predictions that would aid better risk management and provide evidence to inform their decisions.

1.2 *Purpose of risk analysis*

Our paper describes how we can use the empirical data about supplier characteristics and performance records typically held in enterprise resource systems to better understand and predict aspects of supplier risk. These data have been primarily recorded by the manufacturing companies to schedule production and manage operations. We show how such data might also be exploited to provide information to better manage risk in the supply process.

Our goal is to develop a modelling suite that supports analysis of risk in different stages of supplier life so that managers with different levels of responsibilities can use the information appropriately. For example, the information gained from analysis could be used by supply managers: to predict late deliver-

ies given loading and other factors affecting the supplier ability to deliver on time; to provide a robust comparative assessment of supplier performance for variables of interest, such as late delivery rates or non-conformance rates on past and ongoing projects; and to estimate how much it is worth investing in suppliers, new or established, through auditing and other activities to learn more about their capabilities to improve; predictions of the chance of being better off exiting a supplier compared with investing in their development to improve performance.

1.3 Distinctiveness of approach

Our modelling and analysis is distinctive in relation to the growing literature on supply chain risk management. See, e.g. Sohdi and Tang (2012). This is because we ground our models in the theory and methods of stochastic processes, statistical inference and decision analysis to develop sound ways of making best use of typical operational data to directly address routine decisions of supply managers. In contrast, the wider literature on supply chain risk management can be classified as considering, for example: identification of the sources and classes of risk; forensic analysis of major historical supply risk events; mathematical modelling of problem abstractions under assumptions that facilitate modelling but are not necessarily easily capable of instantiation with data in practice; and the use of simulation, both discrete event and system dynamics, to model scenarios where the efficacy of mitigation and controls can be assessed in terms of potential risk reduction.

Our work relates to the wider literature on supply chain analytics and risk analysis. For example, Gallien et al. (2015) report the design and implementation of a modelling system also designed to make use of enterprise resource planning data but with a different purpose – that of optimising production planning. Methodologically our models relate to methods for analysis of data generated from counting processes, e.g. Quigley et al. (2007).

In the remainder of this paper we shall describe our methodology, present examples of analysis using de-sensitised real data based on a study with one manufacturing company, and discuss the implications for managing supply risk.

2 METHODOLOGICAL FRAMEWORK

2.1 Overview

Figure 1 illustrates our methodological process by articulating the relations between the data inputs, the modelling components, and the primary management decisions. Table 1 describes wider classes of decisions associated with each model as well as summarising the underpinning scientific methods. Figure 1 shows how models can be combined to support dif-

ferent types of decision and/or to use information gained from an earlier modelling stage. The three models within the framework are described in the following sub-sections.

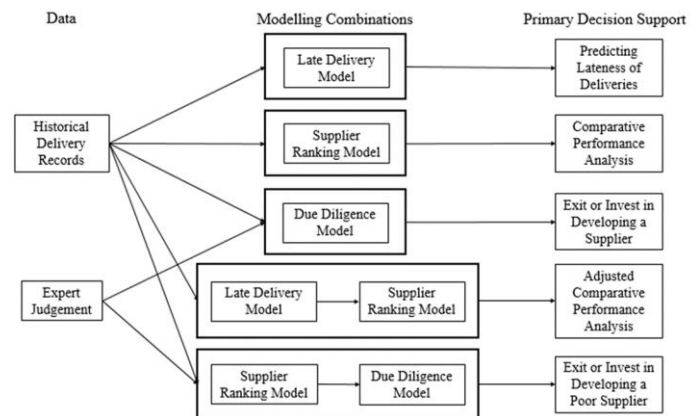


Figure 1. Methodological framework for supply risk analysis.

Table 1. Summary of modelling methods and decisions.

<i>Model</i>	<i>Method</i>	<i>Decision Support</i>
Late Delivery Model	Negative Binomial Regression	<ol style="list-style-type: none"> 1. Predict delivery performance. 2. Identify which orders are of high risk of being delivered late. 3. Support negotiations of batch sizes and lead times. 4. Use to correct for influential factors to support improved ranking.
Supplier Ranking Model	Empirical Bayes inference for ranking of events realised from stochastic counting model with epistemic uncertainties assessed through historical data.	<ol style="list-style-type: none"> 1. Identify the position of a supplier within a pool of similar suppliers. 2. Initiate investigations of underperforming suppliers 3. Support discussion in supplier meetings. 4. Select suppliers for possible development investment.
Due Diligence Model	Decision theory for value of information framed as a Bayesian model representation of epistemic and aleatory uncertainty.	<ol style="list-style-type: none"> 1. Identify the value of learning more about a supplier before investing in development. 2. Determine optimal investment in supplier development. 3. Support decisions when choosing between different learning and development activities.

2.2 Late delivery model

Late deliveries can cause severe impacts on production. Consequences can range from minor inconvenience through to late penalty charges from the end customer. Delivery performance is a critical means for judging a supplier, but represents a highly complex relationship. It reflects not only the supplier, but the relationship between the supplier and manufacturer and the manufacturer's ability to effectively manage the relationship. For manufacturers with global supply chains there are numerous variables believed to influence whether a delivery will arrive on time or be late. Consider the potential events which may cause a delivery to arrive late. Early in the relationship between a manufacturer and a supplier there may be issues associated with the interfaces between the data management systems of either parties, or the manufacturer may be delayed in sending drawing specifications and so hold up progress at the supplier. For the supplier, during production of parts, there may be machine failures, raw material problems, unanticipated rework, and scrappage due to newness of parts. During transportation of parts from supplier to manufacturer, a logistics company may fail to pick up the part on time, shipment may be delayed due to weather, delays at customs may occur, and delays at goods received of the company all may prevent items arriving on-time.

For supply chain managers understanding the consequences of loading decisions is of key importance. Being able to formalise the relationships between, for example, production loading and delivery performance of a supplier is regarded as providing a useful tool. Understanding the associated degree of risk for different orders is required to allow supply managers to better formulate orders as well as to anticipate orders which are of high risk of arriving late.

Data for delivery performance is recorded routinely in the manufacturing company and can be obtained from their enterprise planning system. Each record represents one delivery and contains associated covariate information about the order, i.e. the number of parts ordered, lead time of part, supplier, date ordered, and so on. Naturally, for large manufacturers there exists a large amount of records.

Given the nature of the target variable of interest (i.e. count of number of days or weeks a delivery is late), the suitability of count models are clear for this problem. Further, since we have a set of covariates driving lateness, as well as the large empirical data set that is constructed from delivery records, we adopt the well-established Negative Binomial regression modelling (Hilbe 2007) as the specific type of count model to predict lateness of orders. Count data has advantages in that it collapses the count of interest, in the context discussed here this is the lateness of delivery which is measured as a non-negative integer value, and relates this to a set of covariates

(Hilbe 2007). Data for the covariates can also be taken from the enterprise resource planning system.

2.3 Supplier ranking model

Supply managers currently assess the relative performance of sets of suppliers, although this tends to be based upon simple statistical ratios. For example, late delivery rates given by the number of late deliveries relative to the number of orders for a given time window, or the non-conformance rate given by the number of non-conforming parts relative to the number of parts ordered for orders delivered over a given time horizon. However, as mentioned earlier, the patterns of orders and order sizes is highly variable between different suppliers depending upon the nature of their part and also through time depending upon end customer orders. Methods are required that take into consideration this heterogeneity in the exposure of suppliers to risk.

We have developed an empirical Bayes method for ranking Poisson count data with heterogeneous exposure to risk; see Quigley et al. (in prep.a). Empirical Bayes is an established method of statistical inference and has been used in many risk and insurance applications especially where low frequency, or zero frequency events are the norm (e.g. Quigley et al. 2007). In our research, we have developed a novel means of ranking using empirical Bayes inference that allows us to compute point and interval estimates. The latter are particularly useful as they allow managers to distinguish between statistically different performances between suppliers. Interval estimates help us avoid the pitfalls of naïve over-interpretation of the rankings listed in a league table by allowing suppliers who are clearly in a class of their own to be differentiated from those that are indistinguishable in terms of their performance.

Although our empirical Bayes ranking methods have been motivated for estimation based on the rate of events, it is possible to extend the method to compute ranks that account for the consequence of these events too. For example, the production line is restrictive in the number of products which can be manufactured concurrently. The sequential manner of production means that the cost of internal quality failure associated with the identification of non-conforming parts or late delivery of parts which are needed in that stage of assembly will differ. As the assembly moves through the production line it increases in value and so the hurt cost of unusable or unavailable supplied parts becomes higher. Therefore, the hurt cost of events affecting different production stages varies. Ranking can take into account both the frequency of events and the costs of the consequences of these events.

2.4 Due diligence model

As well as taking decisions to maintain or to exit a relationship with an existing supplier with whom the manufacturer has empirical evidence of past performance, choices can be made about how much to invest in developing new suppliers in the interval between their contract and first delivery. This interval may be of the order of several months and in some cases up to two years. For a new supplier evidence of anticipated performance will be based on data collected during the procurement process, including, for example, audits and parts production tests. Hence for all suppliers, new and established, evidence exists that allows uncertainties in performance to be articulated, although the nature of the data used and the degree of uncertainty may vary between suppliers.

We propose a novel modelling approach to support managers to decide the optimal level of investment in supplier development under uncertainty about performance. This model is applicable for all suppliers. Further, we can provide managers with an assessment of the probability of gaining better performance by selecting a new supplier from the market rather than continuing to use an existing supplier. This can aid decisions concerned with exiting current suppliers, whether these choices are motivated by a poor supplier performance record or by pressures to switch to lower cost suppliers for which there will be uncertainty about performance.

We have developed a Poisson-Gamma probability model within a Bayesian framework, allowing us to represent both the epistemic (state-of-knowledge) and aleatory uncertainties in the event rates. Full mathematical details are given in Quigley et al. (in prep.b). Using this model we can obtain estimates to value a supplier development activity and to assess whether it is worth gaining more information, through, for example, further auditing, to reduce epistemic uncertainty. This model can be populated using empirical data records on, for example, non-conformance frequencies to represent the aleatory uncertainty in the form of a Poisson counting process. The epistemic uncertainty in, say, the true non-conformance rate can be represented through the prior distribution. This prior distribution can be estimated as an empirical prior using empirical Bayes inference if a suitable pool of suppliers can be selected and justified allowing their historical performance data to be used. Or, the prior distribution can be specified subjectively using methods of structured expert judgement. Other parameters required to populate the model include an assessment of the effectiveness of a supplier development activity as well as the financial value of lost production due to an unusable or unavailable part. Supply managers can be supported in expressing their judgements about the likely effectiveness of activities and sensitivity analysis can be run to explore values across a range of effectiveness rate values. We have chosen to express

our model results as a function of the unit cost for lost production as this allows managers to insert values for this unit costs and to scale up as appropriate. This also avoids requiring financially sensitive data to be required as a model input.

3 ANALYSIS OF SUPPLIER RISK

3.1 Overview

We now show a selection of analysis conducted using de-sensitised data from a manufacturer that illustrates the application of the modelling methods and discusses the choices made in data preparation. All analysis shown is “real” insofar as the supplier issues, the data and the results are based on an industry-based study. Note that the applications reported have only been undertaken after the methods themselves (e.g. supplier ranking and due diligence models) have been subject to extensive scientific evaluation theoretically and for simulated experiments. The analysis presented is a selection only and hence our examples draw upon different commodity groups and failure events to provide illustration.

3.2 Prediction of late deliveries

The data used are constructed from the delivery records for all suppliers in the commodity for a period of one year. The variables which are available in the enterprise resource planning system are limited but new variables have been defined and constructed using existing records based on discussion with the supply chain managers. For example, covariates include those capturing supplier loading such as order lines per week, number of parts ordered per week, whether order placed within the lead time and the nature of the logistics used to make delivery.

To give a sense of the data, historical performance data on the aspect of delivery failure is recorded individually. For each delivery received at the manufacturer, there is a record of the delivery. Each entry represents a delivery of a type of part. The entry captures standard delivery details, including the supplier, the part type, lead time of the part, delivered date and expected delivery date. Thus, we have records for every delivery made.

Figure 2 shows the fit of an estimated late delivery model applied to one commodity group on a weekly time interval. This is one of several models developed since we have explored how late delivery performance behaves for different commodity groups and under different time delay intervals. Counts of late deliveries have been aggregated at the weekly level because cyclical data patterns are present in the data at a daily level and these patterns are artificial insofar as they are due to data recording practices rather than reflecting supplier performance. For the commodity model five variables have been

included. By comparing the observed count of late deliveries against the predicted count, Figure 2 shows that visually the model describes the actual performance variation well. It is apparent that the model oscillates between over- and under-predicting but there are no obvious discrepancies between predicted and observed distributions.




Figure 2. Predicted and observed weeks late for commodity data set under Negative Binomial regression model.

The model can be used to generate risk profiles for future order scenarios. Figure 3 compares the risk profiles of four order scenarios, labelled A-D. Each corresponds to different loading and logistic settings. The predicted probability distributions of number of late deliveries of different lengths of delay allow us to explore how different input scenarios are realised as risk profiles. Order scenario C decays slowest, showing more uncertainty and a relatively high chance of very late orders as indicated by the thickest right tail. In contrast, order scenario A decays faster showing there is less chance of arriving late and has a lower expected mean number of days late. The results from such modelling can be presented visually as we have shown, but also numerically. Either way, the analysis allows the order characteristics to be examined before placing the order or prior to scheduled delivery so that action can be taken to avoid or at least mitigate the effects of late delivery on production.




Figure 3. Predicted later delivery risk profiles under four order scenarios (A-D).

3.3 Comparative assessment of suppliers

We show analysis for one commodity group comprising 35 suppliers. Empirical data consists of historical delivery and non-conformance performance records which we have translated into counts of the number of late deliveries, numbers of orders, number of non-conformances and number of parts. For illustration we show analysis for overall performance of each supplier within one financial year. The methods of ranking can be applied to counts over any specified time interval of interest to the managers. For example, the equivalent results can be obtained for counts of failure events over weekly, monthly, quarterly or other time intervals and could also be updated on a rolling basis since this type of analysis relates to regular operational management reviews.

Figures 4 through 6 show different visual presentations of the data analysis. Figure 4 shows an extract from a tabular display ranking the suppliers from best to worst performing in terms of late delivery. The empirical Bayes (EB) late delivery rate is given for each supplier together with the raw data. The empirical Bayes delivery rate is a weighted average of the estimated mean late delivery rate of the pool of 35 suppliers and the classical estimator of the late delivery rate for a given supplier. The mean rank and the 95% interval estimates for the mean rank can be computed. The point estimate of the mean rank naturally starts at 1 for the best performing supplier on delivery and ends at 35 for the worst supplier. The range of the interval corresponds to the length of the bar. The width of the interval estimate can be interpreted as a statistical test to distinguish between those suppliers that are in class of their own (i.e. upper - lower = 0) or whether there are sub-sets of suppliers with similar performance and so “equivalent” ranks (i.e. upper - lower limit > 0). For example, in terms of late delivery performance, supplier S4 is the best, and suppliers S5, S18, S6, S10, S22, S2 are second best and statistically performing around the same level given that we adjust for their different exposures to risk. In contrast, supplier S35 is worst, S20 is second worst and S31 is third worst, while S11, S33, S32 are fourth worst.

Figure 5 shows boxplots of the distribution of ranks for the 35 suppliers for both late delivery and non-conformance. This visual shows the distribution of ranks, where the mean rank is denoted by a diamond and the quartiles by the box, with lines out to the minimum and maximum values. The mean ranks are those given numerically in column 4 of Figure 4. The boxplots provide information about the expected ranks as well as the variation in the distribution. The suppliers are given on the horizontal axis according to their supplier code (i.e. S1, S2, ..., S35). For example, consider supplier S4 whose distributions for both late deliveries and non-conformances indicate less variation than with, say, supplier 34 based on relative spreads of the distributions. Com-

parison between the boxplots indicates that supplier S4 performs well on delivery given the mean and median rank position, but is one of the worst five suppliers in terms of non-conformance and so part quality. Hence, we can gain insight into the uncertainty about performance from the spread of the distributions of ranks, we assess the positioning of a supplier relative to peers for a particular performance measure and we can compare performance across multiple measures for a given supplier.

Supplier	No. orders	No. late deliveries	EB rate	Mean rank	Interval range
S4	1233	89	0.074	1	
S5	962	104	0.111	2	
S18	90	8	0.114	3	
S6	450	49	0.114	4	
S10	158	18	0.127	5	
S22	52	5	0.136	6	
S2	1503	190	0.128	7	
S12	447	95	0.215	8	
S7	603	141	0.235	9	
S9	359	87	0.245	10	
...
S11	1717	768	0.446	30	
S33	12	8	0.499	31	
S32	455	210	0.458	32	
S31	43	25	0.527	33	
S20	181	97	0.523	34	
S35	27	23	0.696	35	

Figure 4. Extract of late delivery ranks from best to worst performing suppliers in pool of 35.

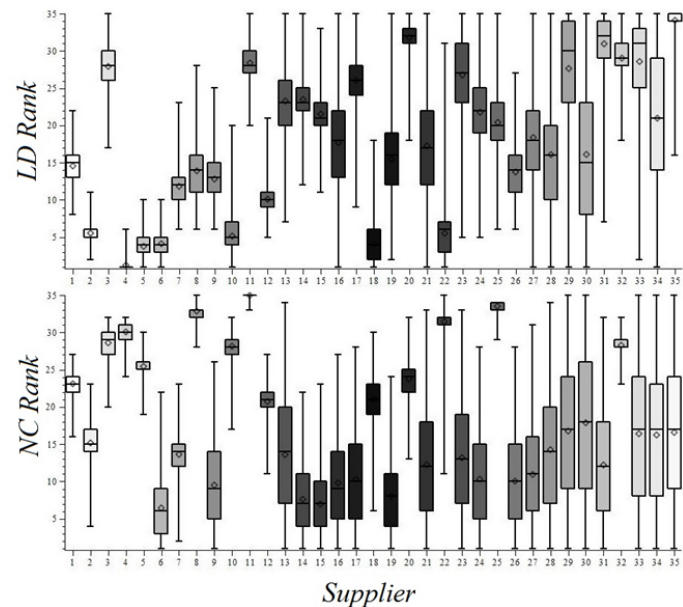


Figure 5. Box plots of distribution of ranks for late deliveries (LD) and non-conformance (NC) for supplier pool.

Figure 6 shows the relationship between the mean rank of late deliveries against the mean rank of non-conformances for the pool of 35 suppliers. The immediate impression is of no clear pattern of relationship given the scatter of points, each representing a

supplier position. However, it is possible to identify the overall best and worst performing suppliers since these will be in the bottom left and top right of the scatter plot given we have defined rank = 1 to be the best. In this data set, supplier S6 is performing well in terms of delivery and quality performance, while the Pareto optimal set of suppliers S11, S32, S20 and S35 are performing poorly on both aspects of performance. Thus the supply managers have information about where to target effort to develop and also where lessons can be learnt from good suppliers through benchmarking.

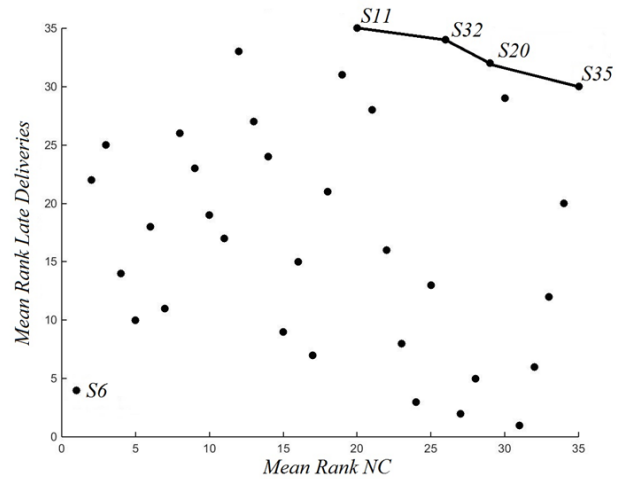


Figure 6. Scatter plot of mean late delivery rank against mean non-conformance rank for pool of 35 suppliers, with Pareto frontiers (top right = worst, bottom left = best).

3.4 Due diligence model

For suppliers identified as performing poorly the decision of whether to exit the supplier or to invest in them through a development activity must be made. Let us show how we may conduct such analysis by focussing on quality performance.

Figure 7 shows the probability of improving performance with respect to non-conformance by replacing a supplier with a random selection from the market for a pool of 200 suppliers across multiple commodity groups. Two clusters of suppliers emerge. Taking 0.5 as natural division, there are a substantial number of suppliers with probabilities less than a 0.5 implying that these are not candidates for exit, but for development. There are around 60 suppliers for which it is more likely that the market will give a better supplier if performance remains at current levels since their probabilities exceed 0.5. A supply manager can determine an appropriate cut-off level with respect to the probability and determine a level which if a supplier is below triggers a flag for possible exit, meaning the supplier would not be contracted for future projects, or for development.

Figure 8 shows the prior distributions of the epistemic uncertainty in the true non-conformance rate

for a supplier of interest and of one randomly chosen from the market and so without any history with the manufacturer. As we gain experience with a supplier, over time the uncertainty will decrease, as depicted with the supplier of interest having much smaller variability. The question is, how much is it worth investing to reduce our epistemic uncertainty about the supplier of interest by conducting developing activities to learn more about their performance capabilities? That is, to reduce the variation of the prior distribution. Also, should we invest to learn more or should we simply decide whether or not to directly invest in development activities?

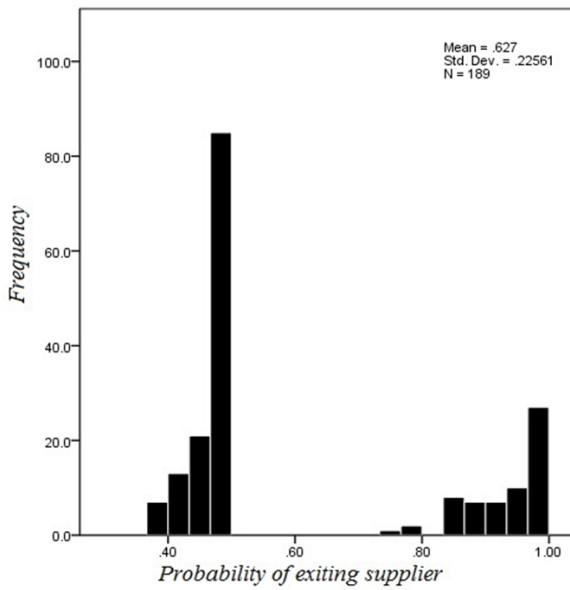


Figure 7. Probability of supplier performing worse than a random selection from a pool.




Figure 8 Prior probability distributions of epistemic uncertainty in true non-conformance rate for supplier of interest (solid) and a supplier from the market (dashed).




Figure 9. Decision box to support assessment of relative added cost-benefit of supplier development under different prior epistemic uncertainty of true non-conformance rate.

Figure 9 shows a simple decision box that provides a visual representation to help answer these questions. The vertical axis represents the prior distribution for the true non-conformance rate and the horizontal axis represents the ratio of the added cost-benefit of a development activity, which is a function of the financial value of lost production, the effectiveness rate of an activity to learn or develop a supplier, and the anticipated reduction (or removal) of epistemic uncertainty due to supplier development. The scaling is such that the 45° line represents the decision rule that distinguishes between situations where it is worth investing and those where investment in supplier development would not be the optimal strategy. For example, in Figure 9, the prior mean of the supplier with the dashed distribution of epistemic uncertainty covers both grey and black zones and as such it may or may not be cost-effective to develop, so would be worth learning more about before deciding on whether to develop. The supplier represented by the solid line falls in the black zone and hence implies that learning more about this supplier will not change the decision, i.e. it is not worth investing in supplier development.

4 IMPLICATIONS AND CONCLUSIONS

The extracts from analysis presented has evolved through a longitudinal study with a manufacturer of bespoke variants of engineering systems assembled using parts sourced from a complex, global supply chain. By engaging over a period of two years, we have had continuous opportunity to learn more about the related issues that concern the supply managers, their desire to have evidence to support and challenge their experience, their relative trust in empirical data and their awareness of the under-utilisation of the information contained in operations databases.

We have adopted an iterative approach to defining model purpose, developing methods with different degrees of scientific novelty, and evaluating applica-

tions within trial contexts using data from multiple commodity sets of suppliers over time horizons spanning several months to several years. To date, relatively more industry feedback has been gained through application of the late delivery prediction model and the supplier performance compared with the due diligence model. All models have been subject to scientific validation, reported elsewhere, but all also require to be subject to further industry validation. In summary, we consider our impact to date to be as follows.

We have shown how empirical records in databases and enterprise resource planning systems originally collected for operations management, can be exploited to provide insights to better understanding of patterns and causal factors to help manage risk.

We have been able to provide statistical evidence through modelling that has been directly relevant to on-going supply management tasks, such as supplier performance ranking and late delivery predictions. The visual forms of evidence presentation has undergone several iterations, and will continue to be refined as we continue to learn about useful ways of presenting mathematical modelling in visual and management friendly ways.

We have scoped a road-map for analysis of supply risk that is grounded in a coherent probability modelling framework and includes different modelling purposes that align with decisions made at different stages in the life of the supplier relationship with the manufacturer. Our models have been populated with variables for which data, empirical and judgemental, are available. Although the general nature of the underlying methods allows for the inclusion of other variables, that are currently acknowledged as possibly influencing supply risk but have not been directly observable so far. We are considering ways of measuring such variables to increase the predictive capability of the late delivery model.

We have explored how forming appropriate pools of suppliers can support better inference and hence assessment of risk. This challenges the conception that suppliers should be pooled according to part type or commodity group only, which while useful for other reasons is not necessarily the most appropriate for probabilistic analysis of risk. Although not shown, we can for example adjust the pools of suppliers to compensate for factors that statistically influence performance through regression modelling.

We have gained a degree of acceptance by supply managers that their state of knowledge (epistemic) uncertainties can be represented in models as prior distributions, although to date we have mainly used empirical priors. Another piece of ongoing work relates to the elicitation of epistemic uncertainties.

We have developed algorithms to support computational implementation of the methods in spreadsheets, which is the default general-purpose analysis tool in the manufacturing companies. Our initial

analysis has been conducted with more specialised software, such as the R package for the Negative Binomial regression analysis and Maple for the supplier ranking and due diligence models. Exact inference can be mathematically challenging and computationally expensive in some instances. Hence research work has been undertaken to create and verify algorithms and approximations to validate them so that our models can be easily applied by practitioners.

Our future work will span industry validation studies of the methods developed and trialled to date and scientific research to extend modelling to include decision models that capture suppliers, as well as the manufacturer, perspectives to inform strategic supply management decisions under uncertainty.

More generally, the methods presented have been explained and used in the context of supply risk management. However, the modelling methods are general and so can be used for risk analysis in other contexts where counts of events require to be modelled to support comparative performance analysis of different entities and estimating the value of investing to improve performance and reduce risk.

REFERENCES

- Hilbe, J.M. 2007. *Negative binomial regression*. Cambridge: Cambridge University Press.
- Gallien, J., Mersereau, A., Novoa, M., Dapena, A., Garro, A. (2015). Initial shipment decisions for new products in Zara. *Operations Research*, Vol. 63.2, pp.269-286.
- Quigley J., Bedford, T., Walls L. (2007). Using empirical Bayes to estimate the rate of occurrence of rare events. *Reliability Engineering and System Safety*, Vol. 92, pp.619-627.
- Quigley, J., Walls, L., Parsa, M. (in prep.a). A ranking method for Poisson count data with heterogeneous exposure: application to operational supply risk management.
- Quigley, J., Walls, L., Demirel, G., Maccarthy, B., Parsa, M. (in prep.b). Is it worth developing a supplier? The value of information under uncertainty. *European Journal of Operational Research*.
- Sohdi, M.S. and Tang, C.S. (2012). *Managing supply chain risk*, Springer.

ACKNOWLEDGEMENTS

We would like to acknowledge the role of the industrial partners who have been critical to the progress of the research. Our research is funded by the UK EPSRC grant (EP/K031686/1) and engagement with the industrial partner has been supported by the EPSRC Impact Accelerator Account.