



Supplier quality improvement: The value of information under uncertainty



John Quigley^a, Lesley Walls^{a,*}, Güven Demirel^{b,c}, Bart L. MacCarthy^b, Mahdi Parsa^a

^a Department of Management Science, University of Strathclyde, William Duncan Building, 130 Rottenrow, Glasgow G4 0GE, UK

^b Operations Management and Information Systems Division, Nottingham University Business School, Jubilee Campus, Nottingham NG8 1BB, UK

^c Management Science and Entrepreneurship Group, Essex Business School, University of Essex, Southend-On-Sea SS1 1LW, UK

ARTICLE INFO

Article history:

Received 8 January 2016

Accepted 21 May 2017

Available online 26 May 2017

Keywords:

Supply chain management

Risk analysis

Uncertainty modelling

Decision analysis

Manufacturing

ABSTRACT

We consider supplier development decisions for prime manufacturers with extensive supply bases producing complex, highly engineered products. We propose a novel modelling approach to support supply chain managers decide the optimal level of investment to improve quality performance under uncertainty. We develop a Poisson–Gamma model within a Bayesian framework, representing both the epistemic and aleatory uncertainties in non-conformance rates. Estimates are obtained to value a supplier quality improvement activity and assess if it is worth gaining more information to reduce epistemic uncertainty. The theoretical properties of our model provide new insights about the relationship between the degree of epistemic uncertainty, the effectiveness of development programmes, and the levels of investment. We find that the optimal level of investment does not have a monotonic relationship with the rate of effectiveness. If investment is deferred until epistemic uncertainty is removed then the expected optimal investment monotonically decreases as prior variance increases but only if the prior mean is above a critical threshold. We develop methods to facilitate practical application of the model to industrial decisions by a) enabling use of the model with typical data available to major companies and b) developing computationally efficient approximations that can be implemented easily. Application to a real industry context illustrates the use of the model to support practical planning decisions to learn more about supplier quality and to invest in improving supplier capability.

© 2017 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY license. (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction and industrial motivation

Our research is motivated by engagement with major manufacturing companies that make complex, high value engineered products. The companies with which we have collaborated are responsible for the design, manufacture and assembly of parts but, given the nature of their final products, are also systems integrators of parts that are procured from global supply chains. The responsibilities of supply chain management within these organisations include selecting and developing suppliers, as well as ensuring a sufficient supply of parts to the required specification to meet production demands. These supply bases are extensive and often there is a long lead time with initial contracting of new suppliers happening 3–5 years ahead of the delivery of supplied parts.

Company operations are underpinned by large databases containing information on suppliers (e.g. commodity grouping, technology maturity, geographical location), items (e.g. unit price, lead time, design ownership), and orders (e.g. volumes, delivery status, quality conformance). Routine management reports include data analysis to provide information about supplier performance. Company cultures encourage and embrace rational analysis for operational decision-making. These include decisions to undertake different kinds of activities for poorly performing suppliers and to plan interactions with some suppliers to avoid future problems. Supplying parts at the required quality level is fundamental to achieve the desired level of performance. Supplier development is a costly activity for the companies because it requires deployment of skilled personnel for substantial periods of time. The deployment of such resources requires consideration of the costs and effectiveness of activities. It is within this industrial context that we seek to help management (1) to assess how much it is worth spending to improve supplier quality performance and (2) to understand whether there is value in learning more about supplier quality capabilities.

* Corresponding author.

E-mail addresses: j.quigley@strath.ac.uk (J. Quigley), lesley.walls@strath.ac.uk (L. Walls), guven.demirel@essex.ac.uk (G. Demirel), bart.maccarthy@nottingham.ac.uk (B.L. MacCarthy), mahdi.parsa@agriculture.gov.au (M. Parsa).

Krause, Handfield, and Scannell (1998) describe supplier development as “any set of activities undertaken by a buying firm to identify, measure and improve supplier performance and facilitate the continuous improvement of the overall value of goods and services supplied to the buying company’s business unit”. In considering the two challenges posed by our industry problem, we distinguish between two types of activity: those that primarily will help us learn more about the state of a supplier’s current capabilities, such as plant visits, auditing (Handley & Gray, 2013; Mayer, Nickerson, & Owan, 2004); and those interventions primarily designed to improve supplier quality, such as supplier training, allocating buyer personnel to improve the supplier’s technical base and operations (Krause, Handfield, & Scannell, 1998; Krause, Handfield, & Tyler, 2007). We can then conceptualise a modelling approach that incorporates a two stage decision process, considering how much should be invested in supplier quality improvement activities and whether it is valuable to make an investment now or after learning more about the supplier. These decisions are made under uncertainty about the true quality level that a supplier will achieve. The degree of uncertainty will be influenced by how much experience the buying firm has with a supplier. For established suppliers with whom the buyer has a long history about quality achieved, the uncertainties may be less than for a supplier who is more recently integrated into the buying firm’s supply base.

To build a meaningful model we need to understand the nature of uncertainties affecting supplier quality performance. Our general model is developed with parameters to reflect quality uncertainties. A distinctive feature of our approach is that we distinguish between aleatory and epistemic uncertainties, which relate respectively to those uncertainties that are regarded as irreducible and those that are reducible if more information is collected (Hoffman & Hammonds, 1994). Generally, epistemic uncertainty represents some degree of ignorance or incomplete information about the system or aspects of the system of interest, and importantly such uncertainty can be reduced as information is collected. In contrast, aleatory uncertainty describes the inherent random variation that is a property of the system and is therefore not considered reducible (Bedford & Cooke, 2001). In operational quality systems an improvement in capability would be realised by a reduction in the process variation resulting from a decision to develop a supplier’s quality performance (Kotz & Lovelace, 1998). Epistemic uncertainty in this context is concerned with the a priori state of knowledge about a supplier’s process capability and is expressed before making the decision to develop a supplier or not. Learning by the buyer about a supplier’s true quality capability reduces epistemic uncertainty.

We develop a stochastic model within a Bayesian framework to capture both the epistemic uncertainty associated with true supplier quality performance as well as the aleatory uncertainty associated with the inherent randomness in a supplier’s performance such as that observed in quality performance data. Our approach is grounded in the value of information concept that data has value if, once analysed, it can result in a change of decision (Ketzenberg, Rosenzweig, Marucheck, & Metters, 2007; Wagner, 1969). We consider value to be a combination of the likelihood of changing a decision and the magnitude of its consequence. By formulating an appropriate stochastic model we can estimate the uncertainty associated with the decision consequences, assign likelihoods to possible data and update the stochastic model in view of data.

We consider a context where we have a dominant prime entity (the buyer), such as one of the major manufacturers with which we work, that relies on an extensive base of suppliers. We therefore assume a single buying organisation with multiple suppliers that have been selected according to the buyer’s standard procurement process. Thus, for a new supplier there is insight into anticipated quality performance based on evidence from, for

example, quality process reviews, quality certification, quality achieved for similar parts, or first article inspections. For those suppliers that have supplied parts to the buyer, data will also exist on quality performance achieved historically. Our model is intended to be most useful for those suppliers whose relationship with the buyer is relatively new and for whom a proactive approach to development will be taken, for example, during the period between signing a contract and delivery of the regular supply of orders. This is because in such cases epistemic uncertainty is likely to be greater than for suppliers with whom the relationship is more mature.

We do not consider the choice of activity beyond the two classes of development noted above; learning and improvement. Our model requires as inputs an expression of the buyer’s assessment of epistemic uncertainty in the true supplier quality, as well as the financial value of production losses that will be incurred by the buyer if sub-standard parts are supplied, and an assessment of the effectiveness of development activity. The model provides the level of the optimal investment in a supplier improvement activity with an upper bound on the amount it is worth spending to reduce the epistemic uncertainty about the supplier quality by targeting learning activities before investing in improvement. Such results help the manager to screen suppliers to assess whether it is worth conducting additional plant visits, audits or other learning activities first, or whether it is more appropriate to invest directly in, for example, training, deployment of buyer resources into the supplier, root cause analysis or other activities aimed at directly making quality improvements.

In this study we address the challenge posed by a practical industry problem by developing and evaluating an innovative and applicable modelling solution using a sound mathematical methodology. Our principal contribution is a new modelling framework for supplier development taking into account the value of information. The model is grounded in the theory of decision analysis and statistical inference, and is aligned with an important industrial supply management problem for which we develop a methodology to support implementation with real data. Our model addresses gaps in the existing literature in relation to research on supplier development and the value of information within a supply chain quality management context. The existing literature tends largely either to develop mathematical models for assumed scenarios providing insightful thinking tools, or to discuss the theory and practice of supplier development in an operational supply management context.

We examine the literature relevant to our problem context and position our work in relation to existing empirical knowledge and models on supplier development in Section 2. Our scientific modelling contribution is described in Section 3. We explain how we formulate the stochastic model based on assumptions about the probabilistic representation of uncertainties and present a number of propositions related to properties of the model. We develop an exact solution for the expected value under perfect information, which is the limiting case of buying down epistemic uncertainty through learning activities. To support practical implementation, we derive a computational approximation and evaluate the conditions under which it is accurate. Section 4 presents an application of our model to real, albeit de-sensitised, industry data on supplier non-conformance rates for a set of key tier 1 suppliers to a large industrial prime. We present an empirical Bayes method to estimate the prior distribution representing the epistemic uncertainty in supplier performance using typical data contained in industry databases. After discussing the reasonableness of our assumptions given the industry problem and data, we present a selection of ways in which the findings of our model can be communicated to supply chain managers. Section 5 presents our conclusions and discusses the implications of our findings for practice and theory,

including identifying future challenges and opportunities for further work.

2. Positioning within scientific literature

Our research relates to several strands of the supply chain management and decision sciences literatures. We briefly review seminal studies on supplier development in order to position our modelling approach appropriately within this context. We critically review those studies that focus on modelling the value of information in supply chains with a particular emphasis on the treatment of uncertainties.

2.1. Supplier development

Krause, Handfield, and Scannell (1998), Che and Hausch (1999), and Krause, Handfield, and Tyler (2007) provide detailed accounts of supplier development approaches in practice. In particular, Krause, Handfield, and Scannell (1998) present a general representation of a supplier development process grounded in an extensive industry survey. From our contemporary company engagement, the identified process still typifies many aspects of current practice. For example, critical commodities and suppliers are identified, key performance areas are targeted, appropriate teams are formed, and activities are selected, implemented and reviewed. Interestingly, one step in the general process notes that “opportunities and probability for improvement” through supplier development should be identified. However, no further consideration is given as to how such probabilities should be expressed, although criteria such as the potential to influence the supplier development process, resources required in terms of people and time, as well as the potential return on investment are discussed. Krause, Handfield, and Scannell (1998) pose the question “what criteria should be used to identify suppliers that have high probability of development success?”. Our model helps to answer this question by estimating the value of gaining more information about supplier quality and providing a probabilistic assessment of the risks of such investments, given the degree of epistemic uncertainty, as well as the buyer’s assessment of the potential to develop the supplier.

Krause, Handfield, and Scannell (1998) classify supplier development activities into reactive and strategic approaches. Reactive approaches are the first stage in the development process where investment is made into poorly performing suppliers to undertake corrective actions. Strategic supplier development, on the other hand, is applied at a more advanced stage where the buyer develops a strategic plan for the supply base to increase the long-term capability of the supply network. We position our approach in between these extremes, essentially as a tuned proactive approach, which estimates the value of collecting further information on supplier capability in order to mitigate the risk of poor quality and avoid extensive exposure to risks of a supplier failing to perform.

Supply chain managers are interested in multiple performance measures; Ward, McCreery, Ritzman, and Sharma (1998) highlight four priorities – quality, delivery, flexibility and cost. Krause, Handfield, and Tyler (2007) note that quality has been recognised as important in manufacturing since the 1980s and continues to be of considerable concern since end customer perceptions of the final product quality will be impacted by the quality of parts manufactured by suppliers. They find that performance outcomes in quality, as well as delivery and flexibility, are affected by direct involvement of the buyer’s personnel in supplier development. Hence, deciding how much to invest in interventions aimed at improving supplier quality remains an important business challenge more generally beyond our motivating industrial problem.

Supplier development has been previously investigated in several modelling studies. Based on the primary methodology used,

we classify the literature into (1) game theoretical studies and (2) stochastic modelling approaches.

Most game theoretical studies focus on strategic supplier development for production cost reduction (Bernstein & Kok, 2009; Iida, 2007; Iyer, Schwarz, & Zenios, 2005; Kim & Netessine, 2013; Qi, Hyun-Soo, & Amitabh, 2015). For instance, Bernstein and Kok (2009) consider cost reduction investments of suppliers in an assembly network where the effectiveness of cost-contingent and target-price contracts in promoting investments and increasing profits is analysed. Similarly, Iida (2007) considers an assembly network where both the buyer and the suppliers might invest in cost reduction, showing that effort compensation and cost sharing agreements can enable supply chain coordination. Although cost reduction effort may be interpreted as a means to satisfy certain quality requirements, quality is not given explicit consideration in these studies. More related to our approach is the study by Zhu, Zhang, and Tsung (2007) that explicitly investigates the improvement of a supplier’s quality where both the buyer and the supplier can invest to decrease the non-conformance rate, showing that investment by only the party with higher investment effectiveness is sufficient unless there are resource constraints. Our research differs from these game theoretical studies in two ways. First, we consider the problem from the buyer’s perspective because we adopt a client decision support focus. Second, our approach is based on real-world data, both empirical and judgemental. In contrast, game-theoretical studies in the literature are more general and make idealistic assumptions in particular regarding uncertainty, as we explore further below.

Stochastic programming has been used to study supplier development in a more limited number of studies (Friedl & Wagner, 2012; Wang, Gilland, & Tomlin, 2010). For instance, Wang, Gilland, and Tomlin (2010) use a two-stage stochastic programming framework where in the first stage the buyer selects the investment levels, and based on their returns, which are subject to variation, the order quantities are selected. Of more interest to our problem, Talluri, Narasimhan, and Chung (2010) and Hosseiniinasab and Ahmadi (2015) study strategic supplier development using Markowitz-type mean-variance risk models to formulate the optimal levels of investment in a set of suppliers. Hosseiniinasab and Ahmadi (2015) note the importance of taking into account future performance and anticipated changes in the development of suppliers. They also discuss the use of databases to identify trends and correlations in supplier performance although they use only synthetically generated data for supplier quality, delivery, price and financial position. Our approach differs from these studies in several ways. First, we consider only quality, unlike authors who focus on multiple performance measures, see, for example, the review by Ho, Xu, and Dey (2010). Second, our model provides the expected return on investment in quality improvement as an output, rather than using it as a model input. Third, as mentioned for game-theoretic models, we use real data rather than synthetically generated data. Fourth, we consider a stochastic modelling framework to account for the potential reduction of epistemic uncertainty, which is not captured in Markowitz-type models.

2.2. Value of information

The established concept of value of information (VOI) in decision analysis is predicated on the ability of additional information to reduce epistemic uncertainty. Since Wagner (1969), much has been written about VOI. In the context of inventory management in particular, the value of sharing information about customer demand, forecasts, inventory level, and production capacity for supply chain coordination, cost reduction, and bullwhip effect mitigation has been widely investigated; see reviews by

Sahin and Robinson (2002) and Ketzenberg, Rosenzweig, Marucheck, and Metters (2007).

In their survey article on inventory management, Ketzenberg, Rosenzweig, Marucheck, and Metters (2007) describe VOI as the marginal improvement in value through additional use of information relative to some base scenario, where the base scenario represents a given set of information that can be compared to the value gained from the so-called information scenario, which is structurally identical to the base scenario except that additional information is shared. The authors argue there is growing interest in VOI because of the increasing opportunities to gain more information due to the growth in e-commerce. They discuss different sources of uncertainty, distinguished as random and systematic, which relate to the stochastic and structural characteristics of the system and so could be considered equivalent to aleatory and epistemic uncertainties, respectively. More generally, much has been written about uncertainty in supply chain management with different classifications being proposed; see, for example, the review by Simangunsong, Hendry, and Stevenson (2012).

Interestingly, Ketzenberg, Rosenzweig, Marucheck, and Metters (2007) also formulate several propositions about VOI in an inventory management context. Of most relevance to us are the following, which we paraphrase as follows: (1) VOI is higher when there is greater uncertainty and (2) VOI is higher when there is increased responsiveness. Based on a regression analysis of the empirical data extracted from their literature review, strong support is found for the second and partial support for the first proposition. In our concluding discussions, we reflect upon these propositions with regard to our modelling theory and application in a quality management context.

The supply chain quality management literature contains articles that focus upon decision models related to supplier quality and include the treatment of uncertainty. In the agency settings of such studies, one or both parties involved in a buyer-supplier relation might benefit from hiding private information, leading to moral hazard and adverse selection problems. In such settings the other party needs to provide incentives to establish coordination or incur an information ‘rent’ to reveal the hidden information. As discussed above in the context of supplier development, Zhu, Zhang, and Tsung (2007) build a model to determine which investment options in quality improvement are optimal for both parties when buyer production is outsourced to a supplier. Aleatory uncertainty in the supplier quality control process is modelled in terms of the non-conformance rate and the quality costs incurred by both the supplier and the buyer are explicated. Although the relative states of knowledge of the buyer and supplier are acknowledged, no consideration is given to the articulation of such epistemic uncertainty as a probability distribution.

In contrast, Lim (2001) and Corbett, Zhou, and Tang (2004) discuss the explicit mathematical representation of uncertainties as prior probabilities in the context of buyer-supplier contracting decisions. Neither study uses the term epistemic uncertainty, but the concept is clear from the explicit consideration given to ex-ante views of buyers and the use of prior probability distributions within the models. Lim (2001) develops a buyer decision model for contract option selections when there is uncertainty in supplier quality; expressing a prior probability on the supplier’s technology type to provide a probabilistic assessment of the fraction of defective parts anticipated to be supplied to the buyer. The increasing role of e-commerce data as a motivation for such modelling is identified, with the authors commenting that the visibility of part quality data afforded by shared database systems can impact the degree of information asymmetry between the buyer and the supplier. This observation is contextually important for our problem.

Corbett, Zhou, and Tang (2004) assume a bilateral buyer-supplier monopoly within which they examine scenarios to assess

VOI of multiple contract types. They assume the supplier holds a prior distribution that expresses her uncertainty about the buyer’s internal variable costs. The decision model is developed for a general prior distribution represented by a continuous probability distribution function, although numerical experiments examine various distributional forms of the assumed prior as a form of sensitivity analysis. Different parameter value sets are selected to investigate the effects of controlling the degree of change in the prior mean and variance. Thus, they are, in effect, exploring the effects of different degrees of epistemic uncertainty on their decisions. We adopt an equivalent approach, although we explore sensitivity to changes in the degree of epistemic uncertainty expressed using real data.

While different ways of mathematically representing prior probabilities have been articulated by Lim (2001) and especially by Corbett, Zhou, and Tang (2004), there has been no consideration of how such distributions might be specified in an industrial decision-making context. We show how typically available industry data can be used to form meaningful, rather than assumed, prior distributions to represent epistemic uncertainty.

3. A modelling framework for valuing supplier development

Our modelling concept is illustrated using a decision tree shown in Fig. 1. The buyer needs to choose whether or not to invest in activities to improve supplier quality (upper two branches) or whether to gather more information to learn about supplier quality capability before investing in improvement activities (lower branch). The decision tree is a visual simplification with a binary (good or poor) representation of supplier quality. Our full model considers the occurrence of poor quality events that risk delaying or disrupting supply to the buyer as measured by the number of non-conformances within some period of buyer exposure to risk. The exposure to risk could be measured by, for example, the calendar time or the number of parts ordered from the supplier. We attach a probability distribution to the uncertainties associated with supplier quality. Not shown in the diagram are the buyer valuations associated with each decision pathway, which we measure as the buyer’s loss due to poor supplier quality. The model allows us to determine the highest amount the buyer be prepared to spend in the time window between contracting the new supplier and delivery of orders to learn more about a supplier to reduce epistemic uncertainty about the true quality performance. Hence it supports the manager in assessing whether it is worth learning more before making choices about improvement activities, or whether it is better to make improvement decisions in light of the current state of knowledge. If the latter option is deemed more worthy, then the model further allows the manager to decide whether to invest, or not, in improvement activities and also how much it is worth spending to improve supplier quality should this option be chosen.

3.1. Modelling assumptions about the stochastic nature of uncertainty

Let N denote the random variable, number of non-conformances, and let t denote the exposure to risk of non-conformances to the buyer. We assume that the mean number of non-conformances is proportional to the exposure to risk. We model N as a Poisson random variable with parameter Λ , which denotes the non-conformance rate in proportion to exposure. Exposure may be measured on a continuous (e.g. time) or discrete (e.g. order size) scale. When exposure is measured by a continuous metric then a Homogeneous Poisson Process (HPP) is rather than as a special case of our model. The Poisson model is the simplest model for the number of non-conformances that is both popular in the literature (Montgomery, 2013) and is reasonable for our empirical data, as we shall show in Section 4. The aleatory uncertainty




Fig. 1. Model concept as a simplified decision tree of buyer decisions (□) and supplier quality uncertainties (○).

representing the natural variation in the non-conformance rate is modelled by the Poisson probability distribution conditioned on knowing $\Lambda = \lambda$ as in Eq. (3.1).

$$P(N = n | \Lambda = \lambda) = \frac{(\lambda t)^n e^{-\lambda t}}{n!}, t > 0, \lambda > 0, n = 0, 1, 2, \dots \quad (3.1)$$

Practically, we shall not precisely know a supplier's true non-conformance rate, and so we describe our epistemic uncertainty on the value of this parameter, λ , through a prior probability distribution, denoted by $\pi(\lambda)$. We assume $\pi(\lambda)$ can be described by a Gamma distribution as shown in Eq. (3.2). The Gamma distribution is a conjugate prior to the Poisson model and hence is mathematically tractable giving it popularity as the Poisson–Gamma model (Carlin & Louis, 2009). More importantly, the Gamma distribution with shape and scale parameters, α and β , respectively, provides a flexible family of distributional shapes through which epistemic uncertainty can be expressed probabilistically.

$$\pi(\lambda) = \frac{\beta^\alpha \lambda^{\alpha-1} e^{-\beta\lambda}}{\Gamma(\alpha)}, \alpha > 0, \beta > 0, \lambda > 0. \quad (3.2)$$

If new data becomes available in the form of n_0 observed non-conformance events for a risk exposure of t_0 then, using Bayes Theorem, we can update the prior in Eq. (3.2) with the new data to obtain the posterior distribution. This will also be in the form of a Gamma distribution with a change in parameters as shown in Eq. (3.3).

$$\pi(\lambda | n_0, t_0) = \frac{(\beta + t_0)^\alpha \lambda^{\alpha+n_0-1} e^{-(\beta+t_0)\lambda}}{\Gamma(\alpha + n_0)}, \alpha > 0, \beta > 0, \lambda > 0, t_0 > 0, n_0 = 0, 1, 2, 3, \dots \quad (3.3)$$

The prior in Eq. (3.2) represents our epistemic uncertainty about the true supplier quality given our knowledge of that supplier to date. Our future experience with the supplier will result in a new prior, which is effectively the posterior distribution expressed in Eq. (3.3). Intuitively, as we gain more knowledge of

the supplier, our epistemic uncertainty should reduce. Asymptotically as t_0 approaches ∞ , our epistemic uncertainty approaches zero, because we shall have learnt everything about the true non-conformance rate and so we are left only with the natural random variation described in Eq. (3.1).

3.2. Specification of a prior distribution

As noted above, the prior distribution represents the buyer's epistemic uncertainty about the true non-conformance rate of an individual supplier of interest. If several suppliers are candidates for development by the buyer, then each supplier will be modelled individually. Although we assume the prior comes from a common family of Gamma distributions, the buyer's epistemic uncertainty about each individual supplier is represented by specifying appropriate parameter values for that supplier.

Alternative modelling tactics can be adopted to specify a prior distribution for a supplier. A prior can be constructed empirically. For example, Quigley and Walls (2017) describe a general structured process for eliciting the domain knowledge of an expert, such as a supply chain manager, to define reference factors upon which relevant empirical data from existing suppliers can be matched to the supplier of interest and subsequently verified as an expression of the epistemic uncertainty in that supplier's true non-conformance rate. In this situation, the parameters of the Gamma prior distribution require to be estimated statistically. We show how this is accomplished using standard approaches in the context of our industry example. Alternatively, the subjective judgement of the buyer can be elicited using a structured expert judgement process to express the buyer's epistemic uncertainty in the supplier true non-conformance rate (O'Hagan et al., 2006; Quigley, Bedford, & Walls, 2008). A subjective prior distribution is appropriate if the expert believes s/he has more information than is contained in the relevant empirical data.

In the subsequent steps of our modelling framework, a prior distribution is treated in the same way regardless of whether it has been constructed empirically or subjectively.

3.3. Optimal investment in supplier quality improvement activity

Let x denote the level of investment in supplier quality improvement, where the anticipated benefit is better capability and reduced non-conformance. Let ν denote the loss incurred by the buyer from a single non-conformance, i.e. the unit cost of internal quality failure to the buyer. Here, we consider ν to be an exogenous parameter that needs to be estimated by the buyer. Porter and Rayner (1992) and Schiffauerova and Thomson (2006) provide general reviews on the costs of quality. Following Nandakumar, Datar, and Akella (1993), in quantifying ν consideration should be given to penalties for production delays due to unavailable or unsuitable parts, inventory holding costs for other parts used in lieu of non-conforming parts, costs for rescheduling and switch-over to other orders, and demand side costs of loss of goodwill, such as customer defection and loss of potential future customers, due to the delays, and so on.

To associate the supplier investment level with the supplier performance, we define γ as a measure of the effectiveness rate of the improvement activity, where higher effectiveness is reflected in larger values of γ . We consider a diminishing marginal return of investment in the reduction of non-conformances. We employ the mathematical formulation used by Porteus (1986) and Zhu, Zhang, and Tsung (2007), where the non-conformance rate reduces at a fraction that decreases exponentially with increasing level of investment. Namely, if the number of non-conformances is N in the absence of a development investment, it is expected to decrease to $Ne^{-\gamma x}$ when x monetary units are invested and the effectiveness rate is γ . The effectiveness of a development activity will depend upon, for example, the type and nature of improvement considered, the familiarity of the buyer with the range of parts supplied and production technologies used, and the nature of the relationship between the buyer and the supplier in addition to the commitment of the particular supplier; these need to be reflected in the value chosen for the parameter γ . In this paper we focus upon modelling the epistemic uncertainty in the non-conformance rate because we wish to apply our model using operational data available to the buyer. The effectiveness rate parameter is represented as a single value, although as we show in our industry example, the sensitivity of results to changes in the specified rate of effectiveness can be examined. Future extensions of our univariate stochastic model could accommodate modelling of epistemic uncertainties on multiple parameters at the cost of increased model complexity, computational and elicitation burdens.

Let P denote the profit function expressed as the difference between the value associated with a reduction in the number of non-conformances through the improvement activity, and the investment level, x , required to undertake the supplier quality improvement, as shown in Eq. (3.4). The initial term represents the reduction in buyer loss due to non-conformances before (νN) and after ($\nu Ne^{-\gamma x}$) quality improvement.

$$P = \nu N(1 - e^{-\gamma x}) - x. \tag{3.4}$$

To obtain the optimal level of investment, we evaluate the expected profit when epistemic uncertainty is represented by the Gamma prior distribution in Eq. (3.2). That is, we are considering the prior information scenario associated with the top two branches of Fig. 1. The expected profit function derived is given in Eq. (3.5).

$$E[P] = \nu \frac{\alpha}{\beta} t (1 - e^{-\gamma x}) - x. \tag{3.5}$$

Table 1
Interpretation of elasticity (ϵ) in terms of effectiveness rate (γ) and optimal investment (x^*).

Elasticity	Impact of percentage increase in γ
$\epsilon \leq -1$	Decrease in x^* of at least a percent
$-1 < \epsilon < 0$	Decrease in x^* less than a percent
$\epsilon = 0$	No change in x^*
$0 < \epsilon < 1$	Increase in x^* less than a percent
$\epsilon = 1$	Increase in x^* of a percent
$\epsilon > 1$	Increase in x^* of greater than a percent

The product of parameters νt in the expected profit given in Eq. (3.5) measures the exposure of the buyer to the benefit of the investment. Consistent with the formulation of the Poisson model, t could be measured by, for example, the number of parts ordered from the supplier or the duration of projects for which it is anticipated that the supplier will work with the buyer, and the cost to the buyer of each non-conforming part is ν . This parametric formulation can also accommodate a Net Present Value (NPV) weighting of future benefits as we show in Appendix A. The optimal investment level, x^* , of the expected profit function is given in Eq. (3.6).

$$x^* = \max \left(0, \frac{\ln(\nu t \gamma) + \ln\left(\frac{\alpha}{\beta}\right)}{\gamma} \right). \tag{3.6}$$

From Eq. (3.6) we can make several observations. First, sufficiently low levels of effectiveness will result in zero investment in supplier improvement activity (middle branch of Fig. 1). Second, and less obvious, the optimal investment level does not have a monotonic relationship with the effectiveness rate. This leads us to formulate Proposition 1. The proof is given in Appendix B.

Proposition 1. *The elasticity, denoted by ϵ , of optimal investment (x^*) with respect to the effectiveness rate (γ)—the ratio of the percentage change in x^* with respect to the percentage change in γ —can be expressed as:*

$$\epsilon = \frac{1}{\ln(\nu t \gamma) + \ln\left(\frac{\alpha}{\beta}\right)} - 1.$$

Implying that if:

$$\ln(\nu t \gamma) + \ln\left(\frac{\alpha}{\beta}\right) > 1,$$

then an increase in the effectiveness rate will result in a decrease in optimal investment.

Table 1 summarises the interpretation of elasticity. Expressions for the expected profit at x^* can be obtained through substitution of Eq. (3.6) into Eq. (3.5) to obtain Eq. (3.7). Note that Eq. (3.7) provides an expectation, whereas the actual future outcome will vary as illustrated in Fig. 1.

$$E[P; x^*] = \begin{cases} \frac{\nu t \gamma \frac{\alpha}{\beta} - 1 - \ln\left(\nu t \gamma \frac{\alpha}{\beta}\right)}{\gamma} & \text{if } \nu t \gamma \frac{\alpha}{\beta} > 1 \\ 0 & \text{if } \nu t \gamma \frac{\alpha}{\beta} \leq 1 \end{cases} \tag{3.7}$$

3.4. Assessing worth of learning before investing based on expected value of perfect information

To provide the buyer with a useful means of assessing whether there is value in activities to learn more about supplier quality we compute the expected profit under an assumption of perfect information. Expected value under perfect information (EVPI) does not indicate how much should be invested in a particular quality improvement investment, which was described in Section 3.3. Rather,

estimating EVPI guides managers on how much it is worth spending to buy down epistemic uncertainty about the supplier's true non-conformance rate before investing in an improvement activity. Specifically, computing the expected value of information as the difference between the EVPI and the expected profit without perfect information provides an assessment of how much it is worth spending, at most, to remove all epistemic uncertainty, and hence provides an upper bound on the amount it would cost to reduce uncertainty if information gained was partial and imperfect. This captures the lower branch of Fig. 1.

For the supplier's true non-conformance rate, Λ , we can determine the optimal investment decision under perfect information, which we denote by X_{pi} as it is a function of the random variable Λ . Eq. (3.8) provides an expression for the expected value of profit under perfect information given our modelling assumptions stated in Section 3.1.

$$E[P|x^* = X_{pi}] = \begin{cases} \frac{vt\gamma\Lambda - \ln(vt\gamma\Lambda)}{\gamma} & \text{if } vt\gamma\Lambda > 1 \\ 0 & \text{if } vt\gamma\Lambda \leq 1 \end{cases} \quad (3.8)$$

Since the true supplier non-conformance rate is not known we take the expectation of Eq. (3.8) with respect to Λ using the prior distribution given in Eq. (3.2). Proposition 2 gives an analytic expression for the EVPI. The proof is shown in Appendix B.

Proposition 2. For the Poisson probability distribution given in Eq. (3.1) with a Gamma prior distribution for true non-conformance rate given in Eq. (3.2) and the objective function of form shown in Eq. (3.5), then the Expected Value under Perfect Information (EVPI) can be expressed as shown in Eq. (3.9).

$$EVPI = \frac{vt\gamma\frac{\alpha}{\beta}\left(1 - F\left(1; \alpha + 1, \frac{\beta}{vt\gamma}\right)\right) - 1 + F\left(1; \alpha, \frac{\beta}{vt\gamma}\right)}{\gamma} + \frac{\sum_{i=1}^{\infty} \sum_{j=0}^i \frac{(i-1)!}{(i-j)!j!} (-1)^j \left(\frac{vt\gamma}{\beta}\right)^j \frac{\Gamma(\alpha+j)}{\Gamma(\alpha)} \left(1 - F\left(1; \alpha + j, \frac{\beta}{vt\gamma}\right)\right)}{\gamma} \quad (3.9)$$

where $F\left(1; \alpha + j, \frac{\beta}{vt\gamma}\right)$ is the cumulative distribution function of a Gamma distribution evaluated at 1 with shape parameter $(\alpha + j)$ and scale parameter $\frac{\beta}{vt\gamma}$ given by

$$F\left(1; \alpha + j, \frac{\beta}{vt\gamma}\right) = \frac{\left(\frac{\beta}{vt\gamma}\right)^{\alpha+j} e^{-\frac{\beta}{vt\gamma}}}{\Gamma(\alpha + j)} \sum_{k=0}^{\infty} \frac{\left(\frac{\beta}{vt\gamma}\right)^k}{\prod_{k_1=0}^k (\alpha + j + k_1)}.$$

Computing the expected value of perfect information (i.e. $EVPI - E[P; x^*]$) allows us to obtain an upper bound on how much it is worth spending to learn more about a supplier before investing in quality improvement. If this difference is less than the expected cost of obtaining the supplier information, then Eq. (3.6) can be used to support the buyer decision to invest, or not, in supplier improvement (first or second branch from top in Fig. 1). Otherwise, the buyer obtains more information first to buy down epistemic uncertainty by learning more about supplier quality (lowest branch in Fig. 1).

3.5. Sensitivity of optimal investment to prior variance

It is interesting to explore how the optimal investment in supplier quality improvement under perfect information, X_{pi} , responds to changes in the degree of epistemic uncertainty. We use the prior standard deviation as a summary measure of epistemic uncertainty. Note also that re-expressing the parameters of the prior distribution in terms of the mean and standard deviation can be useful when communicating results to managers since they are more understandable.

Theorem 1 below shows that for situations where the true non-conformance rate is above the investment threshold, that is $\Lambda > \frac{1}{vt\gamma}$, the mean optimal level of investment under perfect information is a monotonically decreasing function of the epistemic uncertainty associated with the non-conformance rate. The proof is shown in Appendix C. We note that the proof does not require the prior to have the form of a Gamma distribution.

Theorem 1. If the non-conformance rate is greater than the minimum investment threshold, i.e. $\Lambda > \frac{1}{vt\gamma}$, then for a fixed mean non-conformance rate $\mu_{\Lambda > \frac{1}{vt\gamma}} = E\left[\Lambda \mid \Lambda > \frac{1}{vt\gamma}\right]$, the expected optimal investment under perfect information is monotonically decreasing with respect to non-conformance uncertainty, i.e. $\sigma^2_{\Lambda > \frac{1}{vt\gamma}} =$

$$Var\left[\Lambda \mid \Lambda > \frac{1}{vt\gamma}\right]. \text{ Specifically, } \frac{\partial E\left[X_{pi} \mid \Lambda > \frac{1}{vt\gamma}\right]}{\partial \sigma^2_{\Lambda > \frac{1}{vt\gamma}}} < 0.$$

3.6. Approximation for EVPI and computational accuracy

Proposition 2 provides an expression for EVPI in terms of a cumulative Gamma distribution function. However, calculating the EVPI using Eq. (3.9) requires a degree of programming knowledge, which might hinder the practical use of the method. Hence **Proposition 3** below gives an upper bound approximation for the EVPI to facilitate easier application in, for example, spreadsheets. We can also obtain a bound on the error between the true EVPI and its upper bound, as shown in **Proposition 4**, and thus obtain a lower bound on the EVPI. Proofs to both propositions are shown in Appendix B.

Proposition 3. The following expression provides an upper bound (UB) for the EVPI expressed in Eq. (3.9):

$$EVPI \leq vt\frac{\alpha}{\beta} - \frac{\Psi(\alpha) - \ln\left(\frac{\beta}{vt\gamma}\right)}{\gamma} - \left(1 - F\left(1; \alpha, \frac{\beta}{vt\gamma}\right)\right) \quad (3.10)$$

where Ψ is the digamma function.

Proposition 4. The error between the upper bound on the EVPI (UB) (Eq. (3.10)) and the actual EVPI (Eq. (3.9)) can be bounded as follows:

$$\begin{aligned} & vt\frac{\alpha}{\beta} - \frac{\Psi(\alpha) - \ln\left(\frac{\beta}{vt\gamma}\right)}{\gamma} - \left(1 - F\left(1; \alpha, \frac{\beta}{vt\gamma}\right)\right) - EVPI \\ & \leq \frac{\left(\frac{\beta}{vt\gamma}\right)^{\alpha} \ln(\alpha) \left(1 - e^{-\frac{\beta}{vt\gamma}}\right) \alpha^2 (\alpha + 2) + e^{-\frac{\beta}{vt\gamma}} \alpha^3 + \left(2 + e^{-\frac{\beta}{vt\gamma}}\right) \alpha^2 + 3\alpha + 2}{\Gamma(\alpha)\gamma \alpha^2 (\alpha^2 + 3\alpha + 2)}. \end{aligned} \quad (3.11)$$

We can characterise the parameter regions where the UB is a good approximation for the EVPI. That is, where the right-hand-side of Eq. (3.11) is sufficiently small. Note that the bound in Eq. (3.11) is not a monotonic function of α . **Corollary 1** establishes the limits of this bound for either α or β when the other is held fixed, showing that the EVPI converges to the UB in these limits for large α or small β .

Corollary 1. For the limits of the error in the upper bound with respect to α and β , the shape and scale parameters of the Gamma prior distribution respectively, are zero for large α and small β . i.e.:

$$\lim_{\beta \rightarrow 0} \frac{\left(\frac{\beta}{vt\gamma}\right)^{\alpha} \ln(\alpha) \left(1 - e^{-\frac{\beta}{vt\gamma}}\right) \alpha^2 (\alpha + 2) + e^{-\frac{\beta}{vt\gamma}} \alpha^3 + \left(2 + e^{-\frac{\beta}{vt\gamma}}\right) \alpha^2 + 3\alpha + 2}{\Gamma(\alpha)\gamma \alpha^2 (\alpha^2 + 3\alpha + 2)} = 0$$

$$\lim_{\alpha \rightarrow \infty} \frac{\left(\frac{\beta}{vt\gamma}\right)^{\alpha} \ln(\alpha) \left(1 - e^{-\frac{\beta}{vt\gamma}}\right) \alpha^2 (\alpha + 2) + e^{-\frac{\beta}{vt\gamma}} \alpha^3 + \left(2 + e^{-\frac{\beta}{vt\gamma}}\right) \alpha^2 + 3\alpha + 2}{\Gamma(\alpha)\gamma \alpha^2 (\alpha^2 + 3\alpha + 2)} = 0$$




Fig. 2. Process showing relation between inputs and outputs when modelling a new supplier. Dotted lines show an alternative path when modelling an existing supplier.

The accuracy of the UB as an approximation for the actual EVPI is assessed in Appendix C for a range of parameter values for α and β . Our results show that the accuracy of the UB increases as the shape parameter α increases, which is consistent with Corollary 1. In addition, its accuracy also increases for increasing values of the effectiveness rate γ and decreasing values of the scale parameter β .

4. Industry example

Fig. 2 summarises our general modelling framework and we now discuss its application to a real industry problem for a large manufacturing company making highly engineered heavy machinery. The company has an extensive in-bound supply base. Lead times can be long for new projects since initial contracting decisions with critical suppliers can be made several years ahead. During the time period before parts arrive, the company faces the ‘buyer’s dilemma’ addressed by our modelling framework. The core decision problems are whether or not to invest in activities to improve quality or whether to invest in activities to learn more about a supplier’s quality. This dilemma is particularly acute for suppliers that are newly integrated into the buying company’s supply base and for whom there may be little empirical evidence about the required part quality since initial contracting and procurement information is limited to checks on certification, quality processes, and previous quality outcomes for related products.

4.1. Setting model parameters

A modelling choice needs to be made about the approach adopted to specify the prior distribution expressing the buyer’s uncertainty of the new supplier’s true non-conformance rate. In this study, the manager is able to construct a suitable comparator pool of existing suppliers based on reference factors elicited

following the methodology of Quigley and Walls (2017). Moreover, the records taken from the company ERP system provide relevant non-conformance data for the suppliers in the comparator pool. Therefore, we elect to construct an empirical prior distribution in this case. The steps in estimating the empirical prior distribution are described in Section 4.2.

The model also requires as inputs an assessment of the effectiveness of the improvement activity and an estimate of the loss incurred by the manufacturer if non-conforming parts are supplied. In this study, we investigate the impact of setting different effectiveness rates on decision-making to cover a range of degrees of effectiveness for different types of improvement activities, the buyer’s familiarity with the supplier, the parts and technologies used, and different levels of supplier engagement. For the purposes of this example, we set the buyer loss to be one unit per non-conformance occurrence, i.e. $v = 1$, to de-sensitise the cost valuations.

4.2. Estimating an empirical prior distribution

We use empirical data from company databases on the annual frequencies of non-conforming parts recorded over several years for a comparator pool of 35 suppliers. To estimate the parameters of the prior distribution using the selected data we adopt a method, known as empirical Bayes, which has been used in a similar manner in technical risk analysis (Quigley, Bedford, & Walls, 2007; Quigley, Hardman, Bedford, & Walls, 2011) and also more generally (Carlin & Louis, 2009).

In order to obtain the Maximum Likelihood Estimates (MLE) of the prior distribution’s parameters, we require an expression for the predictive distribution that explains the relationship between the observed data and the prior parameters (Good, 1976). Eq. (4.1) shows the predictive distribution as a Negative Binomial distribution, where N_i denotes the number of non-conforming parts for supplier i and t_i represents the exposure to risk for sup-




Fig. 3. (a) Joint 95% confidence region for the empirical prior mean and standard deviation of the true number of non-conforming parts per annum. (b) P–P plot of the predictive distribution against the empirical distribution.

plier i , which in this example is 1 year. We justify this choice of exposure because we want to relate the supply risk to the buyer’s manufacturing projects and so consider the implications when parts may not be available for assembly of the engineering product. The remaining notation is the same as in Section 3.

$$\begin{aligned}
 P(N_i = n_i) &= \int_0^\infty \frac{(\lambda_i t_i)^{n_i} e^{-\lambda_i t_i}}{n_i!} \frac{\beta^\alpha \lambda_i^{\alpha-1} e^{-\beta \lambda_i}}{\Gamma(\alpha)} d\lambda \\
 &= \frac{\Gamma(n_i + \alpha)}{\Gamma(\alpha) n_i!} \left(\frac{\beta}{\beta + t_i} \right)^\alpha \left(\frac{t_i}{\beta + t_i} \right)^{n_i}, \\
 &\quad \alpha > 0, \beta > 0, n_i = 0, 1, 2, \dots
 \end{aligned}
 \tag{4.1}$$

Using the predictive distribution we construct the log-likelihood function for the data across the pool of 35 suppliers as follows:

$$l(\alpha, \beta) \propto \sum_{i=1}^{35} \left[d_i + \alpha \ln \left(\frac{\beta}{\beta + t_i} \right) + n_i \left(\frac{t_i}{\beta + t_i} \right) \right]$$

$$\text{where } d_i = \begin{cases} \sum_{j=0}^{n_i-1} \ln(j + \alpha), & \text{if } n_i \geq 1 \\ 0 & \text{if } n_i = 0. \end{cases}$$

We obtain the MLE of the parameters of the empirical prior to be $\hat{\alpha} = 0.0879$ and $\hat{\beta} = 0.0018$. Construction of joint confidence regions for the prior parameters are obtained using likelihood theory (Lawless, 2003). Here, we re-parameterise the prior parameters to express them in terms of the pool mean non-conformance rate (μ) and standard deviation (σ), which we find are more directly understandable in reasoning about the meaning of the empirical prior distribution with supply chain managers. Fig. 3(a) illustrates the 95% joint confidence region showing strong association between the prior mean and standard deviation, which are 48.83 and 164.71 for this data set, respectively. The mean number of non-conforming parts per annum in the pool is not likely to exceed 400 and the standard deviation in the non-conformance frequency is not likely to be greater than 1400.

Now that we have estimated the prior distribution, we can assess the validity of our modelling assumptions. Fig. 3(b) shows a P–P plot to assess the fit of the estimated predictive distribution model to the empirical data. There is a good fit in both extremes of the distribution, although there are values below the 45 degree reference line in the centre indicating that the model is more con-

servative than the data in this region. However the fit is good in the upper right hand tail, which is important in our risk analysis since this is the high consequence situation where the true supplier non-conformance rate may be relatively high.

We now use the empirical prior distribution as a probability model representing the epistemic uncertainty in the true non-conformance rate of the new supplier of interest. This Gamma prior probability distribution function is shown in Fig. 4(a) and indicates a high probability the true non-conformance rate will be low, but the right tail in the distribution implies there remains a relatively lower chance the true non-conformance rate of the new supplier will be high. Using the MLE, $\hat{\alpha} = 0.0879$ and $\hat{\beta} = 0.0018$, we can update the uncertainty associated with the true non-conformance rate to obtain the predictive distribution for the new supplier in the form of a Negative Binomial distribution with parameters $(\hat{\alpha} + n_i, \hat{\beta} + t_i)$. Since we only have information from assessments obtained at initial contracting for the new supplier, we have no data on the number of non-conformances (i.e. $n_i = t_i = 0$). Fig. 4(b) shows this predictive distribution for the number of non-conforming parts per annum conditional on the occurrence of at least one such event. The conditional distribution allows us to illustrate the thick tail of the distribution which would otherwise be dominated by the outcome of zero non-conformances since this probability is estimated to be 0.57 for this data set. The decay of the tail of this conditional distribution is slow, implying that there is a significant risk of many non-conforming parts being delivered by the new supplier given our current state of knowledge about quality obtained from the pool.

4.3. Optimal investment in supplier quality improvement

So, how much should the company be willing to invest to improve the quality performance of the new supplier given prior levels of epistemic uncertainty?

Using Eq. (3.6) we find that the optimal investment to improve the quality of the new supplier is $15.85v$ when the effectiveness rate of an improvement activity is $\gamma = 0.1$, meaning that we would expect to invest up to nearly sixteen times the buyer loss of a non-conformance in improving supplier non-conformance rate. If the effectiveness rate of an activity is $\gamma = 0.5$, then the optimal value of investment decreases to $6.40v$. Consistent with Proposition 1, we find that optimal investment in improvement activities of the




Fig. 4. (a) Prior distribution of uncertainty in the true non-conformance rate and (b) predictive distribution of number of non-conforming parts conditional on at least one non-conformance for the new supplier.

new supplier is lower for higher effectiveness rates. Note that the expected profits are 22.97ν and 40.44ν , respectively.

However, these are expected profits. Above we noted the pattern of variation shown in the prior distribution. If the true non-conformance rate of the new supplier is low (i.e. realised from the left hand tail of the empirical prior distribution) then there remains a risk that a loss will be incurred by implementing the improvement activity. For example, if the true non-conformance rate is $\lambda < 0.05$ for an effectiveness rate of $\gamma = 0.1$ then the probability of making a loss is 0.51; whereas if $\lambda < 0.15$ for $\gamma = 0.5$ then the probability of making a loss is 0.46. We highlight these insights because they allow managers to appreciate the level of risk associated with making an immediate investment in supplier improvement given prior uncertainty.

We can further examine the relationships between the effectiveness rate, optimal investment and expected profit, given the prior epistemic uncertainty as illustrated in Fig. 5. Fig. 5(a) shows the optimal investment and expected profit profiles, which are both zero until the effectiveness rate is above the investment threshold. Beyond this point, the expected profit increases monotonically with the effectiveness rate at a diminishing marginal rate of increase. Expected investment is highest at low effectiveness rates then decreases as the effectiveness rate increases, implying that the higher effectiveness requires less investment to improve profits. In Fig. 5(a) the expected profit function is constrained to be zero for low values of the effectiveness rate, unlike the surface plot shown in Fig. 5(b) where the zones of expected loss and profit can be identified. When effectiveness rate and optimal investment increase, the expected profit is highest. However, as optimal investment and/or effectiveness decrease, so too does expected profit, with high investment and low effectiveness resulting in expected losses.

4.4. Value of learning more about the supplier before investing in improvement

So, should the company invest in activities to learn more about supplier quality to reduce the epistemic uncertainty about the true non-conformance rate?

Table 2

Quantiles and mean of optimal profits distribution under perfect information (in units of ν) for new suppliers to commodity group.

Statistic	Effectiveness of learning activity	
	$\gamma = 0.1$	$\gamma = 0.5$
Quantile	0.5	0
	0.75	0
	0.90	88.71
	0.99	915.39
Mean	43.53	49.29

When the effectiveness rate $\gamma = 0.1$, we find the expected value of perfect information, that is the difference between the EVPI and the expected profit under prior uncertainty, to be 17.05ν . This implies that if the buyer judges it is worth spending up to just over seventeen times the loss incurred by a supplied non-conformance part to remove uncertainty about quality performance then the best decision is to conduct additional learning before investing in an improvement activity. When the effectiveness rate is $\gamma = 0.5$, the expected value of perfect information reduces to 5.44ν .

We can further examine the likely financial consequences of epistemic uncertainty for the true quality performance of the new supplier. Table 2 shows selected quantiles and the mean of the distribution of optimal profits under an assumption of perfect information corresponding to no epistemic uncertainty. The results presented in Table 2 indicate that investment in improvement is not optimal for a large proportion of suppliers new to the commodity group because we find the optimal profit is zero. However it is clear from values of the quantiles, and especially from the relationship of the median to the mean, that this distribution is right skewed implying there is a small chance the new supplier will merit relatively large investment. For example, 1% of such new suppliers to the commodity group would benefit from an investment at least 900 times the value of a non-conforming part.




Fig. 5. Relationship between expected investment and optimal profit, given a range of effectiveness rates for supplier improvement activity. Note in (a) solid is expected profit and dashed is optimal investment and in (b) blue is expected profit and red is expected loss. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4.5. Implications of epistemic uncertainty for decision-making

Through our modelling we aim to support supply chain managers to make informed decisions. Our goal is not to make the optimal decision itself. For this reason, we have presented a selection of results that are both typical of those shared with managers in the manufacturing company and illustrate the behaviour of the model for different inputs. We summarise our key results in relation to the basic modelling concept shown in Fig. 1. Let us consider the situation where the supply chain manager is concerned with an improvement activity which has an effectiveness rate judged to be $\gamma = 0.5$. Our analysis indicates that the manager should be willing to invest up to $5.44v$ to remove the epistemic uncertainty about the true non-conformance rate. If the costs of gathering additional information to learn more about supplier quality exceed $5.44v$ then the supply manager might decide to invest up to $6.40v$ directly in the improvement activity.

So far, our analysis has only considered the degree of epistemic uncertainty at the two extremes: either as estimated by the empirical prior distribution; or totally removed. To further help the manager develop an understanding of the impact of the degree of epistemic uncertainty on optimal investment levels for supplier improvement, we can also examine the impact of changing the prior standard deviation. Fig. 6 shows changes in the expected optimal investment under perfect information as the standard deviation increases from zero through to 200, which is just above the estimated prior standard deviation of 164.71. We also examine four cases of changes in the prior mean around the point estimate of 48.83 to explore part of the confidence region. As previously, we consider effectiveness rates of 0.1 and 0.5. We find that regardless of the effectiveness rate, the expected optimal investment as a function of prior standard deviation is consistently less for lower prior mean. For equivalent prior mean and standard deviation, the expected optimal investment is lower when effectiveness rate is higher.

In our industry example, the managers know the real value of the loss incurred by non-conformance and, based on their procurement knowledge, have informed opinions about the likely effectiveness rate of an improvement activity as well as the associated cost. Our analysis provides them with a means of expressing their

uncertainty about supplier quality evidenced by their data and allows them to investigate options for supplier development and for information seeking activities with an understanding of the inherent risks to inform their decision.

5. Discussion

Our research has been motivated by engagement with industry practice and addresses an important academic topic on the value of information in supplier development. Consequently we believe we have developed a modelling framework that is both useful to supply chain managers and makes a scientific contribution. Our practical motivation has led us to frame a distinctive decision problem where we focus upon the buyer's dilemma of investing in activities to develop supplier quality performance and we aim to make effective and efficient use of available industry data, both empirical and judgemental. Hence we have presented a modelling solution that fills a gap in the literature between the management considerations of the supplier development process and the science of mathematically modelling abstract decision problems using, for example, stochastic programming or game theoretical approaches.

5.1. Conclusions and contributions

Practically, our modelling process has proved valuable to the industry practitioners with whom we have been collaborating since they need to allocate their limited resources to a range of development activities in the context of an extensive number of suppliers. Importantly, by focusing on the expected value of perfect information we help to quickly identify those new suppliers for which there will be no economic benefit in obtaining any further information before the actual improvement investment and those others for which further information is essential. This type of decision aid is critical for prime companies with large supply bases.

Scientifically, our major contributions and insights are as follows. We provide closed form solutions for the optimal level of investment and expected profit under no and perfect information. We establish that the optimal investment level in supplier improvement does not have a monotonic relationship with the effectiveness rate of that activity. Through Theorem 1 we have pro-




Fig. 6. Relationship between the expected optimal investment under perfect information and the prior standard deviation as measure of epistemic uncertainty when effectiveness rate is controlled to be (a) $\gamma = 0.1$ and (b) $\gamma = 0.5$.

vided distribution-free results on the relationship between mean optimal investment levels and uncertainty of the prior distribution. Mean investment level before the learning stage will inform budgeting, as it provides a prediction on the expected amount of investment to a supplier that will be required post due-diligence. We have shown that, *ceteris paribus*, we would reserve the least amount of resources for the suppliers for whom we have the greatest uncertainty. This may seem counter-intuitive as one might expect to reserve extra resources for improvement for those suppliers about which there is greater uncertainty. However, we stress that the budgeting of reserves is made in advance of learning activities and the actual investment in the supplier will be determined only after further information has been collected. We provide an analytical expression for the EVPI that can be used to assess the benefits of learning more about supplier quality processes before investing in improvement activities. We also derive and evaluate an approximation to the EVPI in the form of an upper and lower bound, which supports practical computations within standard software, such as spreadsheets. By creating a visual representation of the relationships between expected profit, effectiveness and optimal investment, and examining the distribution of optimal profit for a given rate of effectiveness, we can communicate the impact of uncertainty on the risks associated with making decisions to managers, as shown in the industry example reported in Section 4.

We reflect upon our insights in relation to the propositions made by Ketzenberg, Rosenzweig, Maruchek, and Metters (2007) mentioned in Section 2. Although formed from an extensive literature review in the context of inventory management, these propositions also express more widely understood characteristics of value of information, hence providing a suitable level at which to consider the implications of our theoretical findings. Our modelling framework is useful in situations when a decision is likely to be sensitive to uncertainty since management support would not be required if (nearly) perfect information exists about the true non-conformance rate since the need to invest, or not, in a supplier would be obvious. Therefore we are consistent with the

proposition that sensitivity of the decision to uncertainty moderates the relationship between the level of uncertainty and VOI. We find, although have not shown, that the expected value of perfect information increases as the prior variance increases, consistent with the proposition that VOI will be larger when uncertainty is greater. Learning activities are intended to reduce the epistemic uncertainty from the prior level, but the rate of reduction will depend on the activity and so vary between activities. Better learning will be achieved when the prior distribution shifts in location towards good or poor quality levels with less spread and this is in line with the proposition that the VOI increases with respect to the level of marginal information. Our approach is predicated upon the view that information has value if it has the potential to change decisions. Our effectiveness rate of a supplier improvement activity essentially provides a mapping from the current to an intended quality performance state of the supplier and so corresponds to a supplier's ability to respond to buyer-led improvement activities given operating constraints.

5.2. Limitations and further work

We have focused upon deriving analytical expressions for the value of perfect information in the context of supplier development investment decisions. This presumes all epistemic uncertainty is removed and so practically the expected value of perfect information only provides the supply chain manager with an upper limit of how much to spend on learning. We can envisage situations where the manager might consider various activities to learn about supplier quality, implying that reduction in epistemic uncertainty about the true non-conformance rate might vary according to the characteristics of different activities. Thus we may obtain more, but not necessarily perfect, information. Our modelling framework can accommodate this situation allowing the manager to assess the levels of uncertainty associated with the non-conformance rate following a learning activity to determine whether it is cost effective. However, we may be required to use simulations to assess situations where partial information is

gained, learning is not perfect and residual epistemic uncertainty remains.

We have focused upon assessing the upper limit of a further reduction in epistemic uncertainty because we were motivated by the challenge posed by the industrial project. Our industry partners were not concerned with the option of delaying key learning activities; instead they were interested in whether or not to implement the learning activities, such as site visits and additional audits during the early phase of supplier relations. Hence our focus has been on obtaining the expected value of perfect information as an upper limit for all further data gathering processes, which would include real options. A real option approach to this problem would be interesting to develop more formally and would be appropriate to modelling partial information, insofar as another branch could be added to the decision tree to represent a delay in the decision to invest and to consider the NPV of the associated costs and benefits to such an alternative.

We have focused upon decisions relating to a supplier newly integrated to a company's supply base in a context where the lead times allow for both learning and improvement activities to be initiated before the regular supply of parts starts. However, a related problem is that of developing existing suppliers with whom the company has a past relationship. Conceptually, our Bayesian stochastic modelling framework supports decisions regarding existing suppliers since it is possible to determine appropriate probability distributions using relevant historical data for the supplier of interest.

We have assumed a Gamma prior distribution. Our choice is aligned with our underlying probability model, which is sufficiently flexible to represent many epistemic uncertainty scenarios. We make the common assumption that non-conformances follow a Poisson distribution. The assumptions support the mathematics of the methods developed and can be validated using standard statistical model checks. However, now that our framework has been articulated, a future challenge is to develop a wider class of probability models that might be suitable to capture different supplier data patterns. This might be especially useful if we extend the set of performance characteristics beyond quality to, for example, late deliveries, or consider situations when there is anticipated improvement in supplier quality as might be expected for start-up companies or new production technologies.

The EVPI can be expanded to assess the value of learning about the effectiveness parameter γ . Assessing the uncertainty about γ may be complicated by the confounding effect of the supplier's willingness to engage in development activities. Additionally, when there is value in knowing the effectiveness of an intervention prior to engagement then learning about both the non-conformance and effectiveness rates is needed to assess the net impact. Developing a bivariate model to simultaneously assess the EVPI for both non-conformance rate and improvement effectiveness would allow the synergies of learning within activities and the dependency between the uncertainties to be analysed. Modelling the epistemic uncertainty in the effectiveness rate within the model also presents additional challenges for elicitation of the prior.

We express the buyer loss due to a non-conforming part supplied as an unknown parameter, which is typical in the literature. For example, *Ketzenberg, Rosenzweig, Marucheck, and Metters (2007)* find that few studies in an inventory management context report total costs of scenarios considered in value of information analysis in the inventory context. We made this modelling choice partly because of the challenge of accessing financial data and estimating such costs accurately, but also because we found that expressing choices relative to this loss is more useful to supply chain managers since it accords with their practice on penalties. There is a need to provide further guidance in the articulation of these costs even if only for applications support, since we know

from our theoretical and empirical work that they will also impact the optimal decision.

Acknowledgements

This research has been supported by EPSRC grant (EP/K031686/1). We also wish to acknowledge the support of our industrial partners for provoking, challenging and engaging in this research.

Appendix A

Let N_i denote the number of non-conformances realised in the i th future time epoch, v_i denote the value of a non-conformance realised in the i th epoch, t_i denote the exposure to risk in the i th epoch, and r denote the discount for one epoch. As before x is the one-off investment made at time 0 and γ is the effectiveness rate. Then we can express the profit expression in Eq. (3.4) as the NPV of profit, P , as follows

$$P = \sum_{v_i} \frac{v_i N_i (1 - e^{-\gamma x})}{(1 + r)^i} - x$$

The expectation of P with respect to N_i is given by

$$E[P] = \sum_{v_i} \frac{v_i \frac{\alpha}{\beta} t_i (1 - e^{-\gamma x})}{(1 + r)^i} - x = \frac{\alpha}{\beta} (1 - e^{-\gamma x}) \sum_{v_i} \frac{v_i t_i}{(1 + r)^i} - x$$

which can be re-expressed as shown below in the form consistent with the expression in Eq. (3.5).

$$E[P] = \frac{\alpha}{\beta} (1 - e^{-\gamma x}) vt - x,$$

$$\text{where } vt = \sum_{v_i} \frac{v_i t_i}{(1+r)^i}.$$

Appendix B

Proof of Proposition 1.

$$x^* \equiv \frac{\ln(vt\gamma) + \ln\left(\frac{\alpha}{\beta}\right)}{\gamma}, \quad \frac{dx^*}{d\gamma} \equiv \frac{1 - \ln(vt\gamma) - \ln\left(\frac{\alpha}{\beta}\right)}{\gamma^2}$$

$$\varepsilon = \frac{\gamma}{x^*} \frac{dx^*}{d\gamma} = \frac{1 - \ln(vt\gamma) - \ln\left(\frac{\alpha}{\beta}\right)}{\ln(vt\gamma) + \ln\left(\frac{\alpha}{\beta}\right)} = \frac{1}{\ln(vt\gamma) + \ln\left(\frac{\alpha}{\beta}\right)} - 1$$

Proof of Proposition 2.

First we establish the following expression which we use of in the derivation of the proof:

$$\begin{aligned} \int_1^\infty z^j \frac{\left(\frac{\beta}{vt\gamma}\right)^\alpha z^{\alpha-1} e^{-z\frac{\beta}{vt\gamma}}}{\Gamma(\alpha)} dz &= \frac{\Gamma(\alpha + j)}{\Gamma(\alpha) \left(\frac{\beta}{vt\gamma}\right)^j} \int_1^\infty \frac{\left(\frac{\beta}{vt\gamma}\right)^{\alpha+j} z^{\alpha+j-1} e^{-z\frac{\beta}{vt\gamma}}}{\Gamma(\alpha + j)} dz \\ &= \frac{\Gamma(\alpha + j)}{\Gamma(\alpha) \left(\frac{\beta}{vt\gamma}\right)^j} \left(1 - F\left(1; \alpha + j, \frac{\beta}{vt\gamma}\right)\right). \end{aligned}$$

where $F\left(1; \alpha + j, \frac{\beta}{vt\gamma}\right)$ is the cumulative distribution function of a Gamma distribution evaluated at 1 with shape parameter $\alpha + j$ and scale parameter $\frac{\beta}{vt\gamma}$.

Note that:

$$F\left(1; \alpha + j, \frac{\beta}{vt\gamma}\right) = \frac{\left(\frac{\beta}{vt\gamma}\right)^{\alpha+j} e^{-\frac{\beta}{vt\gamma}}}{\Gamma(\alpha + j)} \sum_{k=0}^\infty \frac{\left(\frac{\beta}{vt\gamma}\right)^k}{\prod_{k1=0}^k (\alpha + j + k1)}.$$

Table C.3
Ratio of logarithm of upper bound (UB) to true EVPI over simulated parameter ranges.

Effectiveness rate γ	Shape parameter range α	Scale parameter range β	Min $\ln(\frac{UB}{EVPI})$	α value at min	β value at min	Max $\ln(\frac{UB}{EVPI})$	α value at max	β value at max
0.1	[0.001,0.1]	[0.001,0.1]	0.588	0.1	0.001	16.797	0.001	0.1
0.1	[0.001,0.1]	[0.1,1]	7.489	0.1	0.1	26.522	0.01	1
0.1	[0.001,0.1]	[1,10]	21.858	0.1	1	123.074	0.01	10
0.1	[0.1,1]	[0.001,0.1]	0.125	1	0.01	7.489	0.1	0.1
0.1	[0.1,1]	[0.1,1]	2.339	1	0.1	21.858	0.1	1
0.1	[0.1,1]	[1,10]	15.867	1	1	118.364	0.1	10
0.1	[1,10]	[0.001,0.1]	0.009	10	0.01	2.339	1	0.1
0.1	[1,10]	[0.1,1]	0.125	10	0.1	15.867	1	1
0.1	[1,10]	[1,10]	3.86	10	1	110.875	1	10
0.5	[0.001,0.1]	[0.001,0.1]	0.116	0.1	0.001	13.277	0.001	0.1
0.5	[0.001,0.1]	[0.1,1]	3.907	0.1	0.1	14.607	0.01	1
0.5	[0.001,0.1]	[1,10]	9.93	0.1	1	38.418	0.01	10
0.5	[0.1,1]	[0.001,0.1]	0.011	1	0.01	3.907	0.1	0.1
0.5	[0.1,1]	[0.1,1]	0.215	1	0.1	9.93	0.1	1
0.5	[0.1,1]	[1,10]	15.867	1	1	118.364	0.1	10
0.5	[1,10]	[0.001,0.1]	0.001	10	0.01	0.215	1	0.1
0.5	[1,10]	[0.1,1]	0.011	10	0.1	4.508	1	1
0.5	[1,10]	[1,10]	0.186	10	1	27.37	1	10
0.9	[0.001,0.1]	[0.001,0.1]	0.06	0.1	0.001	12.354	0.001	0.1
0.9	[0.001,0.1]	[0.1,1]	2.973	0.1	0.1	12.5	0.01	1
0.9	[0.001,0.1]	[1,10]	7.807	0.1	1	27.916	0.01	10
0.9	[0.1,1]	[0.0,1]	0.001	1	0.01	2.973	0.1	0.1
0.9	[0.1,1]	[0.1,1]	0.039	1	0.1	7.807	0.1	1
0.9	[0.1,1]	[1,10]	15.867	1	1	118.364	0.1	10
0.9	[1,10]	[0.001,0.1]	0	10	0.01	0.039	1	0.1
0.9	[1,10]	[0.1,1]	0.001	10	0.1	2.427	1	1
0.9	[1,10]	[1,10]	0.017	10	1	17.205	1	10

Now we derive the main result. Let $Z = vt\gamma \Lambda$. Since a Gamma random variable is closed under scale transformation we can express $Z \sim \text{Gamma}(\alpha, \frac{\beta}{vt\gamma})$. We seek the following:

$$E_Z[E[P; x_{PI}|Z]] = \int_1^\infty \frac{z-1-\ln(z)}{\gamma} \frac{(\frac{\beta}{vt\gamma})^\alpha z^{\alpha-1} e^{-z\frac{\beta}{vt\gamma}}}{\Gamma(\alpha)} dz$$

$$= \int_1^\infty \frac{(z-1 + \sum_{i=1}^\infty \frac{(-1)^i (z-1)^i}{i}) (\frac{\beta}{vt\gamma})^\alpha z^{\alpha-1} e^{-z\frac{\beta}{vt\gamma}}}{\gamma \Gamma(\alpha)} dz =$$

$$\int_1^\infty \frac{z-1 + \sum_{i=1}^\infty \frac{(-1)^i (z-1)^i}{i}}{\gamma} \frac{(\frac{\beta}{vt\gamma})^\alpha z^{\alpha-1} e^{-z\frac{\beta}{vt\gamma}}}{\Gamma(\alpha)} dz$$

$$= \frac{vt\gamma \frac{\alpha}{\beta} (1 - F(1; \alpha + 1, \frac{\beta}{vt\gamma})) - 1 + F(1; \alpha, \frac{\beta}{vt\gamma})}{\gamma}$$

$$+ \frac{\sum_{i=1}^\infty \sum_{j=0}^i \frac{(i-1)!}{(i-j)!j!} (-1)^j (\frac{vt\gamma}{\beta})^j \frac{\Gamma(\alpha+j)}{\Gamma(\alpha)} (1 - F(1; \alpha + j, \frac{\beta}{vt\gamma}))}{\gamma}$$

which gives an expression for the EVPI.

Proof of Theorem 1.

The optimal level of investment under perfect information is given by $X_{PI} = \max(0, \ln(vt\gamma \Lambda))$. The expectation of X_{PI} when the non-conformance rate is known to exceed the threshold can be expressed as in the following.

$$E[X_{PI} | \Lambda > \frac{1}{vt\gamma}] = \ln(vt\gamma) + E[\ln(\Lambda) | \Lambda > \frac{1}{vt\gamma}]$$

Consider the following Taylor Expansion of $\ln(\Lambda)$ about $\mu_{\Lambda > \frac{1}{vt\gamma}} = E[\Lambda | \Lambda > \frac{1}{vt\gamma}]$, where Λ is a random variable de-

terministic with variance $\sigma_{\Lambda > \frac{1}{vt\gamma}}^2 = E[(\Lambda - \mu)^2 | \Lambda > \frac{1}{vt\gamma}]$.

$$\ln(\Lambda) = \ln(\mu_{\Lambda > \frac{1}{vt\gamma}}) + \frac{(\Lambda - \mu_{\Lambda > \frac{1}{vt\gamma}})}{\mu_{\Lambda > \frac{1}{vt\gamma}}} + o\left(\left(\frac{\Lambda - \mu_{\Lambda > \frac{1}{vt\gamma}}}{\mu_{\Lambda > \frac{1}{vt\gamma}}}\right)^2\right) *$$

We make two observations:

- i) As $\ln(\Lambda)$ is a concave function then we know $\ln(\Lambda) \leq \ln(\mu_{\Lambda > \frac{1}{vt\gamma}}) + \frac{(\Lambda - \mu_{\Lambda > \frac{1}{vt\gamma}})}{\mu_{\Lambda > \frac{1}{vt\gamma}}}$
- ii) Taking the expectation of both sides of (*) results in the following:

$$E\left[\ln(\Lambda) \middle| \Lambda > \frac{1}{vt\gamma}\right] = \ln(\mu_{\Lambda > \frac{1}{vt\gamma}}) + o\left(E\left[\left(\frac{\Lambda - \mu_{\Lambda > \frac{1}{vt\gamma}}}{\mu_{\Lambda > \frac{1}{vt\gamma}}}\right)^2 \middle| \Lambda > \frac{1}{vt\gamma}\right]\right)$$

Re-arranging we have the following bound.

$$\ln(\mu_{\Lambda > \frac{1}{vt\gamma}}) - E\left[\ln(\Lambda) \middle| \Lambda > \frac{1}{vt\gamma}\right] = o\left(\sigma_{\Lambda > \frac{1}{vt\gamma}}^2\right).$$

So from observation i) we know that

$$E\left[X_{PI} \middle| \Lambda > \frac{1}{vt\gamma}\right] \leq \ln(vt\gamma \mu_{\Lambda > \frac{1}{vt\gamma}})$$

and only achieves equality in the deterministic case, i.e. when $\sigma_{\Lambda > \frac{1}{vt\gamma}}^2 = 0$.

From observation ii) we know that the expected investment decreases as uncertainty increases, i.e. as $\sigma_{\Lambda > \frac{1}{vt\gamma}}^2$ increases.

Proof of Proposition 3.

Following the derivation of Proposition 2, we seek to find:

$$\begin{aligned}
 E_Z[E[P; x_{PI}|Z]] &= \int_1^{\infty} \frac{z - 1 - \ln(z)}{\gamma} \frac{\left(\frac{\beta}{vt\gamma}\right)^\alpha z^{\alpha-1} e^{-z\frac{\beta}{vt\gamma}}}{\Gamma(\alpha)} dz \\
 &= \int_1^{\infty} \frac{z - \ln(z)}{\gamma} \frac{\left(\frac{\beta}{vt\gamma}\right)^\alpha z^{\alpha-1} e^{-z\frac{\beta}{vt\gamma}}}{\Gamma(\alpha)} dz - \int_1^{\infty} \frac{\left(\frac{\beta}{vt\gamma}\right)^\alpha z^{\alpha-1} e^{-z\frac{\beta}{vt\gamma}}}{\gamma\Gamma(\alpha)} dz \leq \\
 &\int_0^{\infty} \frac{z - \ln(z)}{\gamma} \frac{\left(\frac{\beta}{vt\gamma}\right)^\alpha z^{\alpha-1} e^{-z\frac{\beta}{vt\gamma}}}{\Gamma(\alpha)} dz - \left(1 - F\left(1; \alpha, \frac{\beta}{vt\gamma}\right)\right) \\
 &= vt \frac{\alpha}{\beta} - \frac{\Psi(\alpha) - \ln\left(\frac{\beta}{vt\gamma}\right)}{\gamma} - \left(1 - F\left(1; \alpha, \frac{\beta}{vt\gamma}\right)\right).
 \end{aligned}$$

Proof of Proposition 4.

The difference between the EVPI and the upper bound provided in Proposition 3 comes from integration over the range [0, ∞) rather than [1, ∞). As such, the error is given by:

$$\begin{aligned}
 \text{error} &= \int_0^1 \frac{z - \ln(z)}{\gamma} \frac{\left(\frac{\beta}{vt\gamma}\right)^\alpha z^{\alpha-1} e^{-z\frac{\beta}{vt\gamma}}}{\Gamma(\alpha)} dz \\
 &\int_0^1 \frac{z - \ln(z)}{\gamma} \frac{\left(\frac{\beta}{vt\gamma}\right)^\alpha z^{\alpha-1} e^{-z\frac{\beta}{vt\gamma}}}{\Gamma(\alpha)} dz \leq \int_0^1 \frac{z - \ln(z)}{\gamma} \frac{\left(\frac{\beta}{vt\gamma}\right)^\alpha z^{\alpha-1} \left(1 - z\left(1 - e^{-\frac{\beta}{vt\gamma}}\right)\right)}{\Gamma(\alpha)} dz \\
 &= \frac{\left(\frac{\beta}{vt\gamma}\right)^\alpha}{\Gamma(\alpha)\gamma} \int_0^1 (z^\alpha - z^{\alpha-1} \ln(z)) \left(1 - z\left(1 - e^{-\frac{\beta}{vt\gamma}}\right)\right) dz \\
 &= \frac{\left(\frac{\beta}{vt\gamma}\right)^\alpha}{\Gamma(\alpha)\gamma} \int_0^1 \left(z^\alpha - z^{\alpha-1} \ln(z) - \left(1 - e^{-\frac{\beta}{vt\gamma}}\right)z^{\alpha+1} + \left(1 - e^{-\frac{\beta}{vt\gamma}}\right)z^\alpha \ln(z)\right) dz \\
 &= \frac{\left(\frac{\beta}{vt\gamma}\right)^\alpha \ln(\alpha) \left(1 - e^{-\frac{\beta}{vt\gamma}}\right) \alpha^2(\alpha+2) + e^{-\frac{\beta}{vt\gamma}} \alpha^3 + \left(2 + e^{-\frac{\beta}{vt\gamma}}\right) \alpha^2 + 3\alpha + 2}{\Gamma(\alpha)\gamma \alpha^2(\alpha^2 + 3\alpha + 2)}.
 \end{aligned}$$

Appendix C

The purpose of our numerical study is to understand the loss of accuracy of the UB as an approximation to the EVPI. Hence, we focus upon the region where the bound is likely to perform poorly, informed by our theoretical results. Setting $v = t = 1$, we simulate combinations of the remaining parameters in the following ranges: $0.001 \leq \alpha \leq 1$; $0.001 \leq \beta \leq 10$; and for $0.1 \leq \gamma \leq 1.0$. Table C.3 summarises the results for effectiveness rates of 0.1, 0.5, and 0.9, and selected 27 partitions of the parameter space for the shape and scale parameters. Reported are the minimum and maximum values of the ratio of the UB to the EVPI in each partition. More than 20,000 simulations have been run in total to calculate these statistics. Lesser (greater) accuracy is implied when the log ratio of the UB to the true EVPI is larger (smaller). By showing the maximums and minimums of the log ratio over the controlled parameter intervals, we gain insight into the best and worst accuracy within each simulation set.

Our results show that the accuracy of the UB increases as the shape parameter α increases, which is consistent with Corollary 1. In addition, the accuracy of the UB also increases for increasing values of effectiveness rate, γ , and decreasing values of the scale parameter β . Therefore, the upper bound becomes a poorer approximation as the effectiveness rate reduces and the values of the scale parameter increases. This implies that for $\alpha > 1$ and $\beta < 0.1$ then the error between the approximation given by the UB and the true EVPI can be as high as a factor of 10.

References

Bedford, T., & Cooke, R. M. (2001). *Probabilistic risk analysis: Foundations and methods*. Cambridge University Press.

Bernstein, F., & Kok, A. G. (2009). Dynamic cost reduction through process improvement in assembly networks. *Management Science*, 55(4), 552–567.

Carlin, B. P., & Louis, T. A. (2009). *Bayesian methods for data analysis*. Chapman and Hall.

Che, Y. K., & Hausch, D. B. (1999). Cooperative investments and the value of contracting. *American Economic Review*, 89(1), 125–147.

Corbett, C. J., Zhou, D., & Tang, C. S. (2004). Designing supply contracts: Contract type and information asymmetry. *Management Science*, 50(4), 550–559.

Friedl, G., & Wagner, S. M. (2012). Supplier development or supplier switching? *International Journal of Production Research*, 50(11), 3066–3079.

Good, I. J. (1976). The Bayesian influence, or how to sweep subjectivism under the carpet. In W. L. Harper, & C. A. Hooker (Eds.), *Foundations of probability theory, statistical inference, and statistical theories of science* (pp. 125–174). Dordrecht: Springer Netherlands.

Handley, S. M., & Gray, J. V. (2013). Inter-organizational quality management: The use of contractual incentives and monitoring mechanisms with outsourced manufacturing. *Productions and Operations Management*, 22(6), 1540–1556.

Ho, W., Xu, X., & Dey, P. K. (2010). Multi-criteria decision making approaches for supplier evaluation and selection: A literature review. *European Journal of Operational Research*, 202, 16–24.

Hoffman, F. O., & Hammonds, J. S. (1994). Propagation of uncertainty in risk assessments: The need to distinguish between uncertainty due to lack of knowledge and uncertainty due to variability. *Risk Analysis*, 14(5), 707–712.

Hosseiniinasab, A., & Ahmadi, A. (2015). Selecting a supplier portfolio with value, development, and risk consideration. *European Journal of Operational Research*, 245, 146–156.

Iida, T. (2007). Coordination of cooperative cost-reduction efforts in a supply chain partnership. *European Journal of Operational Research*, 222, 180–190.

Iyer, A. V., Schwarz, L. B., & Zenios, S. A. (2005). A principal-agent model for product specification and production. *Management Science*, 51(1), 106–119.

Ketzenberg, M. E., Rosenzweig, E. D., Maruchek, A. E., & Metters, R. D. (2007). A framework for the value of information in inventory replenishment. *European Journal of Operational Research*, 182, 1230–1250.

Kim, S. H., & Netessine, S. (2013). Collaborative cost reduction and component procurement under information asymmetry. *Management Science*, 59(1), 189–206.

Kotz, S., & Lovelace, C. R. (1998). *Process capability indices in theory and practice*. London: Arnold.

Krause, D. R., Handfield, R. B., & Scannell, T. V. (1998). An empirical investigation of supplier development: Reactive and strategic processes. *Journal of Operations Management*, 17, 39–58.

Krause, D. R., Handfield, R. B., & Tyler, B. B. (2007). The relationships between supplier development, commitment, social capital accumulation and performance improvement. *Journal of Operations Management*, 25, 528–545.

Lawless, J. F. (2003). *Statistical models and methods for lifetime data*. John Wiley & Sons.

Lim, W. S. (2001). Producer–supplier contracts with incomplete information. *Management Science*, 47, 709–715.

Mayer, K. J., Nickerson, J. A., & Owan, H. (2004). Are supply and plant inspections complements or substitutes? A strategic and operational assessment of inspection practices in biotechnology. *Management Science*, 50(8), 1064–1081.

Montgomery, D. C. (2013). *Statistical quality control*. John Wiley.

Nandakumar, P., Datar, S. M., & Akella, R. (1993). Models for measuring and accounting for cost of conformance quality. *Management Science*, 39(1), 1–16.

O'Hagan, A., Buck, C., Daneshkhah, A., Eiser, J., Garthwaite, P., Jenkinson, D., & Rakow, T. (2006). *Uncertain judgements: Eliciting experts' probabilities*. Wiley.

Porter, L. J., & Rayner, P. (1992). Quality costing for total quality management. *International Journal of Production Economics*, 27, 69–81.

Porteus, E. L. (1986). Optimal lot sizing, process quality improvement and setup cost reduction. *Operations Research*, 34, 137–144.

Qi, A., Hyun-Soo, A., & Amitabh, S. (2015). Investing in a shared supplier. *Production and Operations Management*, 24(10), 1537–1551.

Quigley, J., Bedford, T., & Walls, L. (2007). Estimating rate of occurrence of rare events with empirical Bayes: A railway application. *Reliability Engineering & System Safety*, 92(5), 619–627.

Quigley, J., Bedford, T., & Walls, L. (2008). Eliciting a prior distribution. In F. Ruggeri, R. Kenett, & F. Faltin (Eds.), *Encyclopedia of statistics in quality and reliability*. Wiley.

Quigley, J., Hardman, G., Bedford, T., & Walls, L. (2011). Merging expert and empirical data for rare event frequency estimation: Pool homogenisation for empirical Bayes models. *Reliability Engineering & System Safety*, 96(6), 687–695.

Quigley, J., & Walls, L. (2017). A methodology for constructing subjective probability distributions with data. Elicitation- the science and art of structuring data, Springer international series in operation research & management science.

Sahin, F., & Robinson, E. P. (2002). Flow coordination and information sharing in supply chains: Review, implications, and directions for future research. *Decision Sciences*, 33(4), 505–536.

Schiffauerova, A., & Thomson, V. (2006). A review of research on cost of quality models and best practices. *International Journal of Quality and Reliability Management*, 23(4), 647–669.

- Simangunsong, E., Hendry, L. C., & Stevenson, M. (2012). Supply-chain uncertainty: A review and theoretical foundation for future research. *International Journal of Production Research*, 50(16), 4493–4523.
- Talluri, S., Narasimhan, R., & Chung, W. (2010). Manufacturer cooperation in supplier development under risk. *European Journal of Operational Research*, 207, 165–173.
- Wagner, H. M. (1969). *Principles of operations research*. Eaglewood Cliffs, NJ.: Prentice-Hall Inc.
- Wang, Y., Gilland, W., & Tomlin, B. (2010). Mitigating supply risk: Dual sourcing or process improvement? *Manufacturing & Service Operations Management*, 12(3), 489–510.
- Ward, P. T., McCreery, J. K., Ritzman, L. P., & Sharma, D. (1998). Competitive priorities in operations management. *Decision Sciences*, 29, 1035–1046.
- Zhu, K., Zhang, R. Q., & Tsung, F. (2007). Pushing quality improvement along supply chains. *Management Science*, 53(3), 421–436.