Journal of Operations Management xxx (xxxx) xxx–xxx

Contents lists available at ScienceDirect

Journal of Operations Management

journal homepage: www.elsevier.com/locate/jom

Valuing supply-chain responsiveness under demand jumps

Işık Biçer[a], Verena Hagspiel[b], Suzanne de Treville[c,*]

[a] Erasmus University, Rotterdam School of Management, PO Box 1738, 3000 DR, Rotterdam, the Netherlands
[b] Norwegian University of Science and Technology, Department of Industrial Economics and Technology Management, 7491, Trondheim, Norway
[c] University of Lausanne, Faculty of Economics and Business, Swiss Finance Institute, 1015, Lausanne, Switzerland

ABSTRACT

As the time between the decision about what to produce and the moment when demand is observed (the *decision lead time*) increases, the demand forecast becomes more uncertain. Uncertainty can increase gradually in decision lead time, or can increase as a dramatic change in median demand. Whether the forecast evolves gradually or in *jumps* has important implications for the value of responsiveness, which we model as the cost premium worth paying to reduce the decision lead time (the *justified cost premium*). Demand uncertainty arising from jumps rather than from constant volatility increases the justified cost premium when an average jump increases median demand, but decreases the justified cost premium when an average jump decreases median demand. We fit our model to two data sets, first publicly available demand data from Reebok, then point-of-sale data from a supermarket chain. Finally, we present two special cases of the model, one covering a sudden loss of demand, and the other a one-time adjustment to median demand.

1. Introduction

Postponing an order quantity decision until demand is known—thus reducing the *decision lead time* to zero—eliminates demand-risk exposure.[1] Conversely, demand-risk exposure tends to increase in the decision lead time, resulting in stockouts or overstocks that generate mismatch costs. The ability to postpone the decision about what to order so that the order quantity can be based on better demand information can be conceptualized as a real option (de Treville and Trigeorgis, 2010), and that option's value can be estimated using quantitative-finance methods. Being able to quantify the value of reducing demand-uncertainty exposure that arises from an increase in the decision lead time transforms time into a decision variable.

The first step in estimating option value is to specify the *forecast-evolution process*: how demand uncertainty increases in decision lead time.[2] The simplest case is the random-walk assumption that underlies the Black-Scholes option-pricing model (Black and Scholes, 1973). Each instant that the decision lead time increases, demand uncertainty increases by a minute amount following a geometric Brownian motion. When this constant-instantaneous-volatility process holds, demand is lognormally distributed with volatility increasing in the square root of the decision lead time. This assumption underlies the Cost-Differential Frontier decision tool proposed by de Treville et al. (2014b) that estimates the cost differential that must be offered by a long-lead-time supplier to compensate for the increase in demand-uncertainty exposure resulting from an increase in decision lead time.[3]

In practice, changes in demand may occur suddenly as a change in median demand (*jump*) rather than as an instantaneous increase in volatility. Demand is frequently subject to jumps: The World Economic Forum in its 2012 report on supply-chain risk attributed 44% of supply-chain disruptions to demand shocks (World Economic Forum, 2012).[4] In finance, the limitations of the Black-Scholes model are well known, but the model is generally used as a reasonable approximation (e.g., Bakshi et al., 1997). When the true forecast-evolution process is subject to jumps but the mismatch cost is estimated assuming that all demand uncertainty emerges from a constant-volatility process, how bad is the error? Does the constant-volatility version of the model give a good enough approximation of the mismatch cost for practical purposes, or does the error impact decision making enough to warrant the use of a more complex model?

To address this question we extend the Cost-Differential Frontier decision tool to include jumps following the classic model proposed by

* Corresponding author.
E-mail address: suzanne.detreville@unil.ch (S. de Treville).

[1] A decision lead time of zero means that the production decision can be postponed enough to permit working with firm demand data. Even with a decision lead time of zero, the delivery lead time may well be positive. It need only be short enough that the production can be decided after demand is observed.

[2] For a stock option, decision lead time translates into the time until the stock price is known.

[3] The tool also calculates the cost premium worth paying to reduce decision lead time.

[4] A sudden shift in median demand represents a common type of demand shock, see https://www.investopedia.com/terms/d/demandshock.asp.

https://doi.org/10.1016/j.jom.2018.06.002
Received 4 October 2017; Received in revised form 20 May 2018; Accepted 5 June 2018
0272-6963/

Merton (1976). We use publicly available demand data from Reebok to gain insight into how the choice of model impacts supply-chain decision making. Parsons (2004) studied the cost of demand-risk exposure faced by Reebok in the context of the exclusive license held during the period 2000–2010 to produce replica jerseys with the National Football League (NFL, see also Graves and Parsons, 2005).[5] Available published Reebok data includes the mean and standard deviation of annual demand for replica jerseys for New England Patriots fans; price, cost, and residual value; and a qualitative description of the many types of demand jumps observed by Reebok. Parsons (2004, pp. 74–75) concluded his analysis of Reebok data by proposing that “perhaps the single greatest opportunity for Reebok is to improve its ability to respond to shifts in demand through shorter lead times.” This data is used by Parsons (2004); Graves and Parsons (2005); Parsons and Graves (2005) to demonstrate the value of postponement. Cattani et al. (2008) cited this work as exemplifying the importance of analyzing the value of responsiveness (see also Uppala, 2016). The decision by the authors of the Reebok study to make their data and analysis publicly available made it possible for us to build directly on their work and use our model to extend their analysis.

Our first result is that the impact of jumps on the cost premium worth paying to reduce decision lead time depends on whether a jump is expected to increase or decrease median demand. If a jump is expected to increase median demand, then treating demand uncertainty as though it came from a constant-volatility process results in an understatement of the justified cost premium. If, however, a jump is expected to reduce median demand, then assuming a constant-volatility process will lead to an overstatement of that cost premium. This result arises from how the jump changes the skewness of the marginal demand density. Jumps that are expected to increase median demand will increase skewness as long as they occur relatively rarely.[6] The resulting increase in right-tail weight increases the value of the option to postpone the production commitment. A jump that reduces median demand reduces skewness, making the postponement option less valuable. Managers with whom we have reviewed this result have found it counterintuitive, as they experience more concern about being stuck with excess inventory if a negative jump occurs than about stocking out following a positive jump.

In order to make the analysis as useful as possible to practitioners, we explore two special cases of jumps that are frequently encountered in practice. The first special case models the risk that demand would be completely lost. In the Reebok case this corresponds to a change of team jersey that reduces demand for the old model to zero. We show that adding any reasonable risk of demand loss to a constant-volatility process substantially increases the justified cost premium. The second special case models a one-time update of median demand such as occurs when decision makers obtain early-sales data, which we use to quantify the impact of a possible Super-Bowl win on the cost premium worth paying to reduce decision lead time. These results are not surprising in their direction, but they are striking in their magnitude. When the jumps that everyone knows to exist are explicitly considered in setting the decision lead time, the company is likely to much more aggressively reduce decision lead time.

A question that arose during the research project was whether demand jumps are experienced in supply chains. To address this question, we randomly selected two products from a supermarket chain and analyzed 100 observations of daily demand from point-of-sale data, then counted how many observations had standardized residuals more than three standard deviations from zero. The first product had four such outliers, indicative of demand jumps, and the second had none. We then considered what would happen if we forced a jump model on a product where it seemed like a constant-volatility assumption would suffice. By moving the threshold defining outliers from the usual three down to 2.52 standard deviations, the number of outliers for the second product increased from zero to four. Interestingly, treating these four points as jumps rather than normal variation for the second product substantially increased the cost-premium frontier. Which representation is correct? Our model cannot say. But, the fact that the option value of responsiveness is quite sensitive to when outliers are assumed to represent demand jumps indicates that this is an area that managers should be pondering.

2. Literature review

When we fit a normal or lognormal distribution to demand data, this is consistent with assuming a constant instantaneous-volatility process: The increase in demand uncertainty as the decision lead time increases can be modeled by a Brownian motion. Assuming that demand follows a normal (lognormal) distribution implies an underlying process that follows an arithmetic (geometric) Brownian motion. The standard deviation of demand (log demand) increases with the square root of the decision lead time. Modeling the evolution of a demand forecast as a constant instantaneous-volatility process is not new in supply-chain research. Hausman (1969) demonstrated that a forecast may plausibly evolve according to a geometric Brownian motion. This research formed the basis for the Martingale Model of Forecast Evolution (e.g., Heath and Jackson, 1994; Milner and Kouvelis, 2005). Oh and Özer (2013) emphasized the importance of accurately capturing the forecast-evolution process when deciding about investing in lead-time reduction, giving as an example the case where manufacturer and supplier have asymmetric information. de Treville et al. (2014b) demonstrated the use of a constant-volatility process to transform decision lead time into a decision variable, showing that an apparently compelling cost differential offered by a long-lead-time supplier may result in an increase in mismatch cost that eliminates the cost advantage when volatility is high and the residual value of the item being acquired is low. Including the value of responsiveness into decision making will frequently change the production-location decision. A summary of the calculations that underlie the Cost-Differential Frontier is given in Appendix E. The literature on the value of lead time is summarized in de Treville et al. (2014b) and de Treville et al. (2014a), so we refer readers to those papers. It is generally agreed that lead-time reduction makes companies better able to respond to demand uncertainty, whether from general randomness or demand shocks, and the increased responsiveness reduces the supply-demand mismatch cost (Fisher and Raman, 1996; Iyer and Bergen, 1997; Milner and Kouvelis, 2005; Lutze and Özer, 2008). de Treville et al. (2014a) applied the insights arising from consideration of forecast evolution in decision making in three industrial settings: one with constant instantaneous volatility, one with stochastic instantaneous volatility due to a bullwhip effect, and one with a simple jump-diffusion process in which there was a risk of demand suddenly dropping to zero. The demand forecasts considered in de Treville et al. (2014b) and here are assumed to eventually converge to the true value of demand. They are also assumed to be unbiased, so that the expected value of any given forecast update is zero (this is explained in further detail in de Treville et al., 2014a). de Treville et al. (2014b) considered instantaneous volatility that was both constant and stochastic, showing that stochasticity in the instantaneous volatility increases the value of lead time, especially when comparing longer decision lead times. Thus, demand risk can be a source of profit for a responsive firm. Similar results are emerging in the field of marketing concerning the use of demand clumpiness to increase profit (e.g., Zhang et al., 2014).

The ability to extract value from responsiveness has turned out to be a key piece of the puzzle to policy makers that recognize the importance of manufacturing to the local economy, and that are seeking to identify what kind of manufacturing has the best chance of being competitive in

[5] The license was won by Nike in 2010.

[6] As jumps become less rare, their impact on the marginal density moves from the tail to the body.