Long-Term Challenges in Demand Forecasting and Dissemination

Robyn Louise Goldsmith, MMath (Hons.), M.Res

Submitted for the degree of Doctor of Philosophy
at Lancaster University.

June 2025

Abstract

Accurate demand forecasts are imperative for organisations to make strategic long-term decisions. Demand forecasting for the long-term is challenging and further complications can arise when a demand series has periods where no demand is observed. Motivated by operational decisions made at an automotive manufacturer, we contribute novel methodology to determine long-term forecasts for spare parts.

We first introduce a stochastic demand forecasting model for products in the final phase of the life cycle when demand is in decline. Theoretical results on the bias and variance of the parameter estimates motivate an extension which uses the demand history of parts with the same declining pattern. In experiments on real data, we demonstrate that our extension reduces the mean absolute percentage error, achieves a higher fill rate and incurs less leftover inventory.

We then outline an approach for long-term demand forecasting throughout the product life cycle. We extend our model by pooling the incomplete demand histories of products with similar life cycle behaviour to estimate joint model parameters. We validate our approach on 175 automotive spare parts and find that our extension improves forecast accuracy even for cases when the peak of demand is yet to be observed.

As a third contribution, we develop material to communicate forecasting and modelling topics to wider audiences. We design outreach content based on core principles and consider aims to address recruitment shortages and gender disparity in the mathematical sciences. We reflect on our impact using teacher feedback.

Acknowledgements

I began this project in 2021 under the co-supervision of John Boylan. John sadly passed away in July 2023 after a short illness. It goes without saying that this project would not have been the same without him. I will continue to remember John as a supervisor who had unrestrained enthusiasm for his work and excellent taste in chocolate cake.

To my supervisor Anna, it is impossible to summarise all you have done for me over almost four years in just a few lines. I am so grateful to have been supervised by you. Throughout this journey, I have been awed by your outstanding memory and the way you could finish my sentences when I lost all the right words - somehow you always knew what I was trying to say. Thank you for supporting my ideas, being generous with your time and plainly telling me to book a holiday (more than once). I am very grateful that you were the person alongside me in the circumstances around John's illness and passing. I have really loved working together and would not have hoped to have done this project under anyone else's supervision. Thank you for everything.

I would like to thank the automotive manufacturer that part-sponsored this PhD. Thank you to Mike, Maren, Johnny and the inventory team who provided an interesting real-world problem, exciting data and all their expertise.

I am grateful to have completed this PhD at the STOR-i Centre for Doctoral Training. Thank you to Nicky, Wendy, Kim and Keilah for all you do to support our PhD experiences. I am grateful to the wider leadership team for providing so many opportunities and for encouraging my outreach work. I would like to especially thank Jon

Tawn for all the support you have offered throughout this journey.

My fellow students at STOR-i have provided so many of the highs over these last few years. I want to thank the cohort I started the programme with in 2020 (whilst remaining strictly two metres apart from each other). Thank you to my desk partner, Ziyang, for brightening every tough day in the office and to Ben for lending an ear for all things inventory (and international film) related. I am especially grateful to those I lived with in Lancaster, Rebecca - for our bond over noughties TV shows and feminine rage - and Conor and Owen - who offered much-needed levity and occasionally did my washing up.

I would not have started this programme without the generosity of the lecturers who taught me at the University of Greenwich. I am grateful for the advice and encouragement you gave me (and so many of you continue to give me). A special thank you to my Master's supervisor, Vitaly, for giving me my first taste of research and recommending Lancaster. Thank you also to my undergraduate personal tutor, Erwin, for encouraging me to look for PhD programmes and supporting my application to STOR-i.

I am very grateful to my Mum, Dad and sister, Erika. I did not know anyone who had done a PhD growing up and my decision to do one (in maths of all subjects) must have seemed a bit left-field. I cannot thank you enough for always encouraging me to go for what I want and supporting me in every way you can. Thank you to the rest of my family and my friends for rooting for me from start to finish.

I would also like to thank my best friend and partner, Ross. My pursuit of this PhD has meant that we have spent the last few years living almost 300 miles apart. Thank you for sticking with me, putting things back into perspective in the moments it all became too much and for trying your best to describe my work when your friends say, "So, what is it that Robyn's doing then?"

Thank you to everyone who has been a part of this journey.

Declaration

I declare that the work in this thesis has been done by myself and has not been submitted elsewhere for the award of any other degree.

A version of Chapter 2 has been submitted for publication as Goldsmith, R. L, Sachs, A. L, & Boylan, J. E. (2025). Forecasting Declining Demand in the End-Of-Life Phase.

A version of Chapter 3 has been submitted for publication as Goldsmith, R. L, Sachs, A. L. (2025). Spare Part Demand Forecasting in Every Phase: A Data Pooling Approach to the Bass Life Cycle Model.

A version of the article in Section 4.3.2 has been published as Goldsmith, R (2023). Surfing the Korean Wave. *Significance Magazine* 20(6): 38-41.

A version of the article in Section 4.3.3 has been published as Goldsmith, R (2021). Could Mathematics Be Big at the Box Office? *Mathematics Today* 57(6): 231-233.

Robyn Louise Goldsmith

Contents

List of Figures

List of Tables

List of Abbreviations

EOL	End-of-Life
LTB	Last Time Buy
MPE	Mean Percentage Error
MAPE	Mean Absolute Percentage Error
RMSSE	Root Mean Squared Scaled Error
STEM	Science, Technology, Engineering and Mathematics
OR	Operational Research

Chapter 1

Introduction

Managerial decisions made for the long term have high stakes and include the strategic planning of resources, stock and finances. These decisions are often made under uncertainty, and hence, the decision-makers require forecasts to better anticipate future demand. Long-term forecasts are made for planning horizons of several years and must incorporate long-standing patterns in demand. Unreliable demand forecasts often result in poor decisions, potentially causing significant financial losses or environmental consequences. Although crucial, long-term forecasting is inherently more difficult than near-future forecasting due to the unknown information during the long planning horizon (Simon et al., 2005; Dombi et al., 2018).

Motivated by the strategic decisions faced by an automotive manufacturer, we focus on demand forecasts required for long-term decisions related to the management of spare parts. In this application, inaccurate forecasts may result in inadequate decision-making, causing product shortages, customer dissatisfaction or production of excess stock. While existing approaches focus on short-term decisions (Croston, 1972; Syntetos et al., 2005; Snyder et al., 2012; Sarlo et al., 2023), accurate long-term demand forecasts are essential for strategic planning and generate important benefits both financially and environmentally for a wide range of organisations and application areas. This thesis concerns the role of forecasting and modelling in overcoming long-term challenges.

We first focus on demand forecasting to support long-term operational decision-making for spare parts in the final phase of the life cycle, known as the end-of-life (EOL) phase. This is a relevant task across many industries that operate aftermarkets, including the automotive, aerospace, home appliance, medical, telecommunication and electronic sectors (Durugbo, 2020). Forecasting in the EOL phase is challenging due to a sustained declining pattern in demand. Demand for spare parts can also be very small and intermittent, where in some periods no demand is observed. As a result of the decline in demand, this can be the case even for products that are high volume at the beginning of the EOL phase.

Demand forecasts for spare parts that incorporate the long-term declining pattern are necessary for many important managerial decisions. For example, an inventory decision commonly made for aftermarket items is a final order, also known as a 'last time buy' (LTB), which is intended to cover anticipated demand over a planning horizon of several years. LTB decisions are often prompted by suppliers ceasing the production of parts. This can be a result of technological advancements or, usually if demand has entered the EOL phase, the reluctance of suppliers to fulfil smaller order quantities (Behfard et al., 2018). As a final purchase, LTB decisions are critical. An LTB quantity too low results in unmet service level requirements which can damage a company's reputation. On the other hand, an LTB that is larger than demand produces industrial waste which is harmful to the environment and causes unnecessary production, storage and disposal costs.

A large body of the literature on forecasting the demand for products in the EOL phases has investigated using external factors to model demand (Ritchie and Wilcox, 1977; Hong et al., 2008; Dekker et al., 2013; Chou et al., 2016; Kim et al., 2017; Van der Auweraer and Boute, 2019). Implementation of causal models can bring about new complications for manufacturers in determining relevant explanatory variables. In some cases, relevant information may be spread across different systems or not recorded at all (Syntetos et al., 2016). A different approach is to use historical data to infer the patterns of demand. While some research on time-series models has been undertaken (Moore, 1971; Fortuin, 1980), the literature on time-series approaches to determine long-term forecasts for spare parts in the EOL phase is limited.

As a first contribution of this thesis, we develop forecasting methodology to support long-term decision-making occurring in the latter part of the product life cycle, where demand is in decline. We introduce a stochastic demand forecasting model which incorporates a decline in demand values, in line with products in the end-of-life phase. The model assumes that demand is independently distributed according to the Poisson distribution and that the mean of the demand declines geometrically over time. The model has two parameters, a rate of decline and an initial mean demand size, which corresponds to the mean demand level at the start of the EOL phase and before decline occurs. Maximum likelihood estimation is used to estimate the parameters. We obtain theoretical insights into the bias and variance of the model parameters which indicate how the characteristics of the demand series influence the performance of the model. In particular, we show that more accurate forecasts can be determined for series with higher initial mean demand, corresponding to faster-moving products. This motivates an extension to improve forecast performance across slower-moving products. The model extension pools the demand history from multiple SKUs with the same decline rate to improve the estimation of the model parameters.

We assess the performance of our approach using simulated data across a range of parameter values. Later, we apply our model and its extension to automotive spare parts. We test the efficacy of our approaches with respect to both forecast accuracy and inventory performance. To calculate LTB quantities, we assume a base stock policy and find order quantities for a 95% fill rate using the Poisson loss function. We calculate the percentage error, absolute percentage error, achieved fill rate and percentage of leftover inventory for the total demand over the planning horizon. Overall, our approaches perform well on the spare parts from the automotive manufacturer. We discover that more accurate forecasts are determined for faster-moving products with higher initial demand volumes. On average, we find that forecasts determined using our model extension are more accurate than forecasts determined using our initial approach. Further, forecast accuracy improves as the number of series used to determine the model parameters increases. The results of the experiments echo our theoretical findings.

Our forecasting methodology is of high practical value and has motivated developing a novel decision support system at the automotive manufacturer. The new system is based on the forecast methodology and is currently being piloted to facilitate long-term inventory decision-making. Results from the implementation show that forecast accuracy has improved notably, resulting in less inventory waste, lower purchase and inventory holding costs, as well as increased process efficiency.

The decline in demand that occurs in the EOL phase is not the only pattern that needs to be reflected in forecasts to facilitate effective long-term decision-making. Prior to the final phase, the spare part life cycle is comprised of two other phases which can each last a number of years (Fortuin, 1980; Dekker et al., 2013). In the first phase, demand has an increasing pattern. In the second phase, demand flattens and reaches a peak level. Forecasting demand over the product's life cycle is frequently made difficult due to an absence of historical demand. Demand may have only been observed in the early phases of the life cycle where a declining pattern is yet to be seen and, in some cases, even prior to demand reaching a peak level. Nonetheless, decision-makers may need to determine a forecast that covers the remainder of the product's life.

Observing demand over a fraction of the life cycle can incur inaccurately estimated model parameters which lead to poor forecasts (Srinivasan and Mason, 1986; Van den Bulte and Lilien, 1997; Putsis Jr and Srinfvasan, 2000; Meade and Islam, 2006). Therefore, many current approaches forecast the demand for products with insufficient demand history using the full demand histories of similar discontinued products (Ismail and Abu, 2013; Ganjeizadeh et al., 2017; Dombi et al., 2018; Hu et al., 2019; Li et al., 2021). However, this is not always possible in practice. Spare part life cycles can span over a decade (Koopman, 2011) meaning that current systems may not hold complete records for former products. Another challenge arises when products have periods of zero demand, as is common for spare parts. Forecasting approaches for series with intermittent demand have thus far focused on stationary or declining demand patterns (Moore, 1971; Croston, 1972; Fortuin, 1980; Syntetos and Boylan, 2005).

Our second contribution is methodology to determine long-term forecasts for products throughout the life cycle. We model demand throughout the life cycle as a Poisson process with a non-stationary intensity function. We adopt an established product life cycle curve, the Bass curve (Bass, 1969), to represent the intensity function of the Poisson process. We extend our model using a data pooling approach using the likeness between life cycle patterns of multiple products to improve forecast accuracy. In doing so, we utilise the incomplete demand history of multiple products to determine common model parameters. We test the methodology in a controlled setting to assess the impact of volume size, length of demand history and values of parameters on forecast performance. Thereafter, we conduct a real data experiment using monthly demand information for 175 automotive products. We determine the percentage error and the absolute percentage error for the total demand over the planning horizon to reflect decisions that require a forecast of the aggregated demand over time, such as the LTB. Further, we evaluate performance using a scale-independent metric designed for demand series with zero demand periods, the root mean squared scaled error (Makridakis et al., 2022). Our findings show that our model extension improves the accuracy of forecasts even in cases when the demand history is restricted to the early phases of the life cycle.

The applicability of forecasting and modelling techniques to real-world problems necessitates effective dissemination. Professionals must effectively communicate with stakeholders to ensure the development and implementation of methodology. More widely, operational research (OR) problem settings have relevance to people's everyday lives. A smooth transfer of knowledge helps foster understanding between professionals and the general public. Effectual dissemination of mathematical ideas is also crucial in recruiting future generations of statistical professionals. There is currently a shortage of individuals with the skills necessary to combat challenges of a mathematical and scientific nature facing the UK (Royal Society, 2021; Campaign for Science and Engineering (CaSE), 2023). Moreover, the mathematical community experiences unequal representation of particular demographics. In the UK, women hold only a fraction of roles within science, technology, engineering and mathematics (STEM) sectors (Science and Technology Committee, 2023). In order to ensure the continuation of methodology to solve problems of prosperity related to society, the environment and the economy, the effective promotion of statistical and OR methodology is paramount.

As the final contribution of this thesis, we introduce forecasting and modelling-related content designed for outreach activity. Based on core design principles, we introduce magazine articles and content for outreach talks and workshops aimed at secondary school students. We analyse contributing factors to the shortages and underrepresentation of women in the wider STEM workforce and incorporate themes that address stereotypes related to the competence and identity of professionals within the mathematical sciences. Through a reflection of feedback, we find that the content introduced contributes to the aims of outreach and has a positive impact on student retention. Furthermore, we argue that designing outreach content to target barriers to belonging is impactful toward the aim of achieving equal representation and diversity across mathematical disciplines.

The remainder of this thesis is comprised of four chapters.

Chapter 2 focuses on forecasting challenges related to products in the end-of-life phase of the life cycle. A novel stochastic forecasting model is introduced and theoretical properties of the estimators of the model parameters are derived. A model extension is derived and an evaluation of performance is conducted on both simulated and real data.

In Chapter 3, we consider cases where forecasts may be determined throughout the lifecycle of a product. We introduce a long-term stochastic forecasting methodology for decision-making throughout the life cycle and extend our model using a data pooling approach. Our approach works well using monthly demand data for automotive spare parts.

In Chapter 4, we consider forecasting and modelling in a broader context as a vehicle for engagement and outreach. We present a portfolio of work that is used to communicate forecasting and modelling methodology to lay audiences and discuss the incorporation of themes related to belonging to maximise impact.

Chapter 5 summarises the key contributions of each chapter and discusses possible topics for further research.

Chapter 2

Forecasting Declining Demand in the End-of-Life Phase

2.1 Introduction

Accurate demand forecasts are crucial for making long-term inventory decisions required for parts that have reached the end-of-life phase. However, demand forecasting at the latter phase of the product life cycle is a particularly challenging task because demand for the majority of items is declining. Due to the decline, demand during the end-of-life phase can become very small with periods of zero demand even in cases with high initial values. Standard time-series methods, including those for intermittent demand (Croston, 1972; Syntetos et al., 2005; Snyder et al., 2012; Sarlo et al., 2023), are unbefitting for products in the EOL phase as they are not designed to capture a long-term declining pattern.

Demand forecasts for products with declining demand patterns are relevant for several important business decisions. For example, inventory and warehouse decisions for spare parts in aftermarket industries (e.g., automotive, aerospace and consumer electronics (Durugbo, 2020)) are often regulated by requirements from warranty and service level agreements. Being able to obtain spare parts is also crucial from a customer's point of view and impacts a company's reputation. Therefore, it is of great interest to develop stochastic forecasting models that accurately incorporate the presence of a declining trend expected of EOL products. Inderfurth and Kleber (2013) emphasise the substantial repercussions forecast errors can have for spare part procurement over a long service period.

Companies make so-called 'last time buy' decisions, where they place a last (often very large) order that is supposed to satisfy anticipated demand over a planning horizon spanning many years. If they purchase too little, they are not able to meet service level requirements, which can incur high penalties. If they purchase too much, they have to pay excess production, inventory and holding costs as well as discard the items, which is a waste of natural resources. Consequently, LTB decisions to determine inventory levels and warehouse space requirements for many years ahead are high-stake - both financially and environmentally - and require reliable demand forecasts for the total demand over a planning horizon of several years. Methods for obtaining LTB quantities have been developed by Teunter and Haneveld (1998); Teunter and Fortuin (1999); Van der Heijden and Iskandar (2013); Behfard et al. (2015) and Behfard et al. (2018).

We consider the example of an automotive manufacturer that sells both standard and luxury vehicles. Examples of spare parts in the EOL phase are presented in Figure 2.1.1. Each example stock-keeping unit (SKU) shows demand observed over a nine-year period. Representative of EOL inventories, we show some SKUs for which demand is slower-moving and some which are faster-moving. In these examples, we notice a sustained decline and periods of zero demand. Observing intermittent demands is also in line with the literature on forecasting demand and regular replenishments for spare parts (Van der Auweraer and Boute, 2019; Van der Auweraer et al., 2019; Boylan and Syntetos, 2021).

To model this demand pattern, we propose a forecasting algorithm that estimates

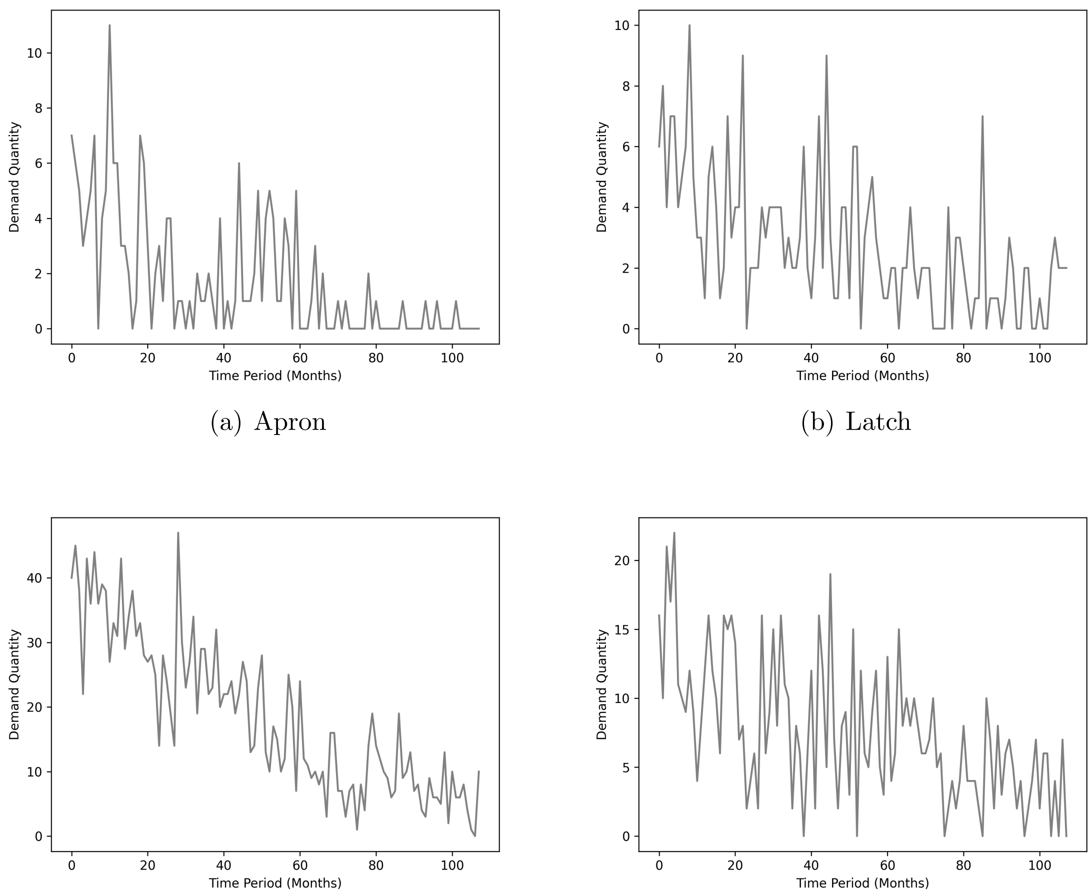

(a) Apron (b) Latch

(c) Brake (d) Shock Absorber

Figure 2.1.1: Demand for SKUs Belonging to an Automotive Manufacturer

the peak demand and the decline rate. We account for the stochastic nature of sales by adopting the representation of demand in each period as an independent Poisson random variable. Trend is incorporated geometrically using a constant decline rate. We suggest that commonality across series of EOL SKUs can be exploited to improve estimates of the parameters, on the assumption that within inventories there exist products which have the same pattern of decline in mean demand size over time.

Our contribution is a stochastic forecasting model for demand that incorporates a decline in demand values befitting of products in the end-of-life phase. We derive theoretical results that show that the performance of the model varies depending on characteristics of the demand series. This motivates a multiple series extension based on the assumption of homogenous decline rates, in which demand observations from other series are utilised to estimate the model parameters. We show numerically that our model extension improves forecast accuracy whereby parameter estimates are obtained using maximum likelihood estimation. Our model performs well on real demand data of spare parts from an automotive manufacturer and our multiple series extension achieves superior forecast accuracy and inventory performance than the benchmark method.

2.2 Literature Review

A large body of the literature on demand forecasting for items in the EOL phase has considered the relationship between exogenous variables and demand. Central to causal models for spare parts, common EOL items, is the idea of the installed base, defined as the number of whole products in use in the population. Dekker et al. (2013) champion employing installed base information to make better demand forecasts for spare parts. Emphasising the intuition behind installed base models, Van der Auweraer et al. (2019) argue for the benefits of incorporating 'real drivers' of demand in forecasting models for practitioners. An early model by Ritchie and Wilcox (1977) utilises concepts from renewal theory, based on the assumption that the decline in the demand for parts is argued to occur due to a lack of customer incentive to replace a failed part as time goes on, opting instead to scrap or replace products. Hong et al. (2008) encapsulate this idea by incorporating a replacement probability into an installed base model. A regression forecasting approach presented by Chou et al. (2016) uses installed base information to estimate demand. Their regression on failure probability provides more accurate forecasts than on their regression on historical sales data. Kim et al. (2017) also present regression-based concepts, incorporating product age as well as the installed base factor. Van der Auweraer and Boute (2019) integrate preventative measures such as scheduled maintenance in their installed base approach.

Other authors have focused on time-series approaches, using demand history to identify patterns and forecast future demand. Moore (1971) is among the first to address demand decline in the latter part of the life cycle. This approach involves the application of deterioration curves to sales data on a logarithmic scale. Fortuin (1980) analyse the LTB problem in the context of consumer electronics. For long-term service periods, they consider that demand follows a non-stationary and uncorrelated Gaussian process with a mean that decays exponentially over time.

Mehringer and Menden (2019) propose a nested workflow of methods, notably clustering, classification and functional approximation to overcome the challenge of forecasting with incomplete historical demand data. The aim is to use similar SKUs, for which more historical data exists, to forecast the demand for the item of interest. However, this clustering and classification framework for intermittent items is still in its infancy.

The seminal approach for forecasting the demand of intermittent items with stationary demand is Croston's method (Croston, 1972), later modified by (Syntetos and Boylan, 2005). Simple exponential smoothing is used to update estimates for non-zero demand sizes and intervals between demand arrivals separately, with the ratio of non-zero demand sizes over intervals giving the demand estimate per period. Others have aimed to tackle the issue of SKUs becoming obsolete over time. Aggregating demand to determine forecasts for products with an increased risk of obsolescence is suggested by Sanguri et al. (2024). Teunter et al. (2011) propose a method, where the forecast declines exponentially in long periods where no demand is observed. A method by Prestwich et al. (2014) propose a forecast that declines hyperbolically in periods of sustained zeroes. Prestwich et al. (2021) adopt linear decay to resolve the issue of a product suddenly becoming obsolete. The decline in these forecasts is incorporated based on no demand being observed. The decline in observed values expected in the EOL phase is not considered in these methods. These approaches are advantageous when regularly monitoring whether or not to replenish stock over concerns of obsolescence. Forecasting procedures for final orders do not fit within this context.

Motivated by the gap in the literature for forecasting spare parts with trend, Altay et al. (2008) utilise the Wright Modified Holt method based on Wright (1986) for forecasting irregularly spaced data. This adaptation aims to fit a linear function to the demand data, using simple exponential smoothing to update slope and intercept estimates. The approach produces only one-step ahead forecasts and is not intended for long-term purchase decisions. Linearity constrains the method's application to demand series with a downward trend, as demand forecasts may result in negative values. Despite intervals between non-zero demands being incorporated into the method in a Croston style approach, estimates are not updated in periods with zero demand in this approach.

The current literature has limitations. Although causal models have been shown to be useful in many cases to forecast demand for spare parts, they can only be implemented if companies can determine relevant information, which can be difficult due to technical restrictions (e.g., some information may be unrecorded or only available in different systems), and may pose additional challenges such as deciding which explanatory variables should be selected (Syntetos et al., 2016). An alternative is to use only data on historical demand patterns. However, the literature on time-series forecasting approaches for EOL products is not extensive. The early work by Fortuin (1980) assumes demand is Gaussian but EOL products can have small demand quantities for which a discrete distribution would be more appropriate. Additionally, the demand level at the start of the EOL phases is specified using data from only one period. The method developed by Moore (1971) is likely to be unsuitable if there is only a small amount of available sales data after the peak of demand is observed. Additionally, the method relies on the representation of demand history on a logarithmic scale rendering it unsuitable for items with zero demand periods. Well-established forecasting methods for intermittent items (Croston, 1972; Syntetos and Boylan, 2005) have thus far overlooked the decline in demand necessary for LTB scenarios. Where trend has been considered (Teunter et al., 2011; Prestwich et al., 2014, 2021), methodology is driven by short-term risks of obsolescence opposed to long-term inventory decision-making.

2.3 Demand Forecasting

In this section, we introduce a stochastic forecasting model for products in the EOL phase. Our model assumes that products are in the declining phase of demand and that demand observations are independent of one another. We represent demand per period as Poisson, following a common representation of spare part demand in the literature (Teunter and Haneveld, 1998; Syntetos et al., 2012; Boylan and Syntetos, 2021). We first present an approach that uses only the previous sales history of a single series to estimate the model parameters. We then extend the model to incorporate the sales history of multiple series to aid the estimation of the parameters for a single series. This is done under the assumption of homogeneous decline rates across series.

2.3.1 Single Series Model

The model we present stipulates that for an SKU in the end-of-life phase, a demand size in period t, Y_t, is an independent Poisson random variable with mean λ_t, where λ_t declines geometrically. We write

$$Y_t \sim \text{ Poisson}(\lambda_t), \text{ where } \lambda_t = \lambda_{t-1}\rho = \lambda_0\rho^t.$$

Here, ρ is a rate of decline and λ_0 is the initial mean when the demand decline begins, such that $0 < \rho < 1$ and $\lambda_0 > 0$. Hence, in period $t = 1$, demand follows a Poisson distribution with mean $\lambda_1 = \rho\lambda_0$. In period $t = 2$ the mean of the Poisson distribution becomes $\lambda_2 = \rho^2\lambda_0$ and so on. This means, that all demand originates in the initial mean λ_0, which is when the peak of the demand occurred and subsequently declines further with every period that has passed.

Estimates from a sample of size n ($\hat{\rho}_n$ and $\hat{\lambda}_{0,n}$) of the parameters (ρ and λ_0) are obtained by applying maximum likelihood estimation on historical demand data. For n independent demand observations, $y_1, ..., y_n$, the likelihood function is

$$\mathcal{L}(\lambda_0, \rho; y_1, ..., y_n) = \prod_{t=1}^{n} \frac{(\lambda_0\rho^t)^{y_t} e^{-(\lambda_0\rho^t)}}{y_t!}$$

and the log-likelihood function is

$$\ell(\lambda_0, \rho; y_1, ..., y_n) = \sum_{t=1}^{n} \left[y_t \ln(\lambda_0\rho^t) - \lambda_0\rho^t - \ln(y_t!)\right]. \tag{2.3.1}$$

The model parameters λ_0 and ρ that maximize the likelihood function over the space $\lambda_0 > 0$ and $0 < \rho < 1$ are given by solving equations (2.3.2) and (2.3.3), respectively,

$$\hat{\lambda}_{0,n} = \frac{(1-\hat{\rho}_n)\sum_{t=1}^{n} y_t}{\hat{\rho}_n(1-\hat{\rho}_n{}^n)} \tag{2.3.2}$$

$$\frac{1}{1-\hat{\rho}_n} - \frac{n\hat{\rho}_n{}^n}{1-\hat{\rho}_n{}^n} - \frac{\sum_{t=1}^{n} t y_t}{\sum_{t=1}^{n} y_t} = 0. \tag{2.3.3}$$

For the derivation of (2.3.2) and (2.3.3), see Appendix 2.A.

2.3.2 Theoretical Properties

Products in EOL inventories can vary greatly in demand volume. In some cases, demand for some SKUs can be very small, posing additional challenges for forecasting. It is therefore important to understand how the performance of models designed for EOL SKUs are affected by demand volume.

As $n \to \infty$, the middle term on the left-hand side of (2.3.3) tends to zero. Hence, a closed form approximation for the maximum likelihood estimator can be determined. Denoting this approximate value by $\hat{\rho}_n^A$ we have

$$\hat{\rho}_n^A = 1 - \frac{\sum_{t=1}^{n} y_t}{\sum_{t=1}^{n} t y_t}$$

and correspondingly the approximate maximum likelihood estimator $\hat{\lambda}_{0,n}^A$ is

$$\hat{\lambda}_{0,n}^A = \frac{(1 - \hat{\rho}_n^A)\sum_{t=1}^{n} y_t}{\hat{\rho}_n^A}.$$

When $n \to \infty$, $\hat{\rho}_n \to \hat{\rho}$ and $\hat{\rho}_n^A \to \hat{\rho}^A$ with both limits $\hat{\rho}$ and $\hat{\rho}^A$ being the same value. Similarly, when $n \to \infty$, $\hat{\lambda}_{0,n} \to \hat{\lambda}_0$ and $\hat{\lambda}_{0,n}^A \to \hat{\lambda}_0^A$ with $\hat{\lambda}_0 = \hat{\lambda}_0^A$. The limit for the estimator for the rate of decline is

$$\hat{\rho} = \hat{\rho}^A = 1 - \frac{\sum_{t=1}^{\infty} y_t}{\sum_{t=1}^{\infty} t y_t} \tag{2.3.4}$$

and the limit for the estimator for the initial mean demand size is

$$\hat{\lambda}_0 = \hat{\lambda}_0^A = \frac{(1 - \hat{\rho})\sum_{t=1}^{\infty} y_t}{\hat{\rho}}. \tag{2.3.5}$$

We derive theoretical insights about the behaviour of the model dependent on the characteristics of the demand series by investigating the bias and variance of the limiting estimators (2.3.4) and (2.3.5). We obtain insights for finite n in the simulation in Section 2.5.

We cannot use the standard asymptotic results of normality for maximum likelihood estimators (Pawitan, 2001) to study the behaviour of the maximum likelihood estimators for this inference problem. This is true even in the limit case when $n \to \infty$. The reasons for this are as follows. We represent demand as non-identically distributed Poisson variables but critically as $n \to \infty$ the mean and the variance of these variables tends to zero and the Fisher Information tends to a finite limit, see Appendix 2.B.

Hence, after some finite value of n all subsequent realisations of the variables will be zero. As a consequence of this property, the key central limit result behind the standard asymptotic normality results will not hold. Crucially, this means that even when $n \to \infty$ we only obtain a finite set of information about ρ and λ_0 and therefore there is no guarantee that the bias of $\hat{\rho}_n$ or $\hat{\rho}_n^A$ or of $\hat{\lambda}_{0,n}$ or $\hat{\lambda}_{0,n}^A$ will tend to zero as $n \to \infty$. Similarly, the variances of these estimators will also not tend to zero as $n \to \infty$.

To study the behaviour of the limiting estimator $\hat{\rho}$ we consider

$$\hat{\rho} = 1 - \frac{\sum_{t=1}^{\infty} Y_t}{\sum_{t=1}^{\infty} tY_t}$$

where $Y_t \sim \text{Poisson}(\lambda_0 \rho^t)$ for all $t \in \mathbb{N}$. We employ approximations for the mean and variance of $\hat{\rho}$ and $\hat{\lambda}_0$ using results on the behaviour of a ratio of random variables using second-order Taylor series expansion (Stuart and Ord, 1998; Elandt-Johnson and Johnson, 1980) to find expressions for the bias and variance of (2.3.4) and (2.3.5). The following propositions show how characteristics of the demand series affect the accuracy of the parameter estimates of a model.

Proposition 1: For independently distributed Poisson demand with a geometrically declining mean, the bias and variance of the limiting maximum likelihood estimate for the rate of decline can be obtained. The second order Taylor series expansion of the expectation and variance of a ratio gives

$$\text{Bias}(\hat{\rho}) = -\frac{(1-\rho)^2}{\lambda_0}, \qquad \text{Var}(\hat{\rho}) = \frac{(1-\rho)^3}{\lambda_0}.$$

The proof of Proposition 1 is given in Appendix 2.C. Proposition 1 leads to the following corollary:

Corollary 1: For a single series with initial mean demand size λ_0 and rate of decline ρ, the $\text{Bias}(\hat{\rho}) \to 0$ and $\text{Var}(\hat{\rho}) \to 0$ as $\lambda_0 \to \infty$.

To study the behaviour of the limiting estimator $\hat{\lambda}_0$ we consider

$$\hat{\lambda}_0 = \frac{(1-\hat{\rho})\sum_{t=1}^{\infty} Y_t}{\hat{\rho}}.$$

where $Y_t \sim \text{Poisson}(\lambda_0 \rho^t)$ for all $t \in \mathbb{N}$.

Proposition 2: For independently distributed Poisson demand with a geometrically declining mean, the bias and variance of the limiting maximum likelihood estimate for the initial mean demand size can be obtained. The second order Taylor series expansion of the expectation and variance of a ratio gives

$$\text{Bias}(\hat{\lambda}_0) = \frac{(1-\rho)^2}{\rho^3\lambda_0} + \frac{1-\rho}{\rho^2}, \qquad \text{Var}(\hat{\lambda}_0) = \varphi_1\left(\lambda_0, \rho\right).$$

The proof of Proposition 2 and the functional form of φ_1 are given in Appendix 2.D.

Corollary 2: For a single series with initial mean demand size λ_0 and rate of decline ρ, the $\text{Bias}(\hat{\lambda}_0)$ and $\text{Var}(\hat{\lambda}_0)$ do not tend to 0 as $\lambda_0 \to \infty$.

We find that the bias and variance of the limiting parameter estimates are functions of the true initial mean demand and decline rate of the demand series. The bias and variance of the decline rate estimate reduces as the true initial mean demand size increases. The variance of estimate for the initial mean demand size increases in line with volume size. The bias of $\hat{\lambda}_0$ decreases as the true initial mean demand size increases. However, a small bias remains for this parameter. The remaining bias, $(1-\rho)/\rho^2$, decreases with respect to the decline rate, whereby slower decline (higher values of ρ) incurs a smaller bias, and tends to 0 as ρ approaches 1. Assuming ρ is not less than 0.9, then the remaining bias is small i.e., less than or equal to 0.13. We note that rates of decline are usually high and a rate of 0.9 is lower than those we observe in the real data example in Section 2.6.

We gain important insights from this analysis. Firstly, we observe that high decline rates reduce the bias of the initial mean demand size estimate. Therefore, demand series with slow decline patterns will produce better estimates of this parameter than demand series that decline rapidly. More notably, the bias and variance of the approximate decline rate estimate and the bias of the approximate initial mean demand size estimate reduce as the true initial mean demand size increases. Consequently, forecasts will be more accurate for demand series with larger initial mean volume sizes and less accurate for demand series of low volume.

The theoretical properties of our model show that the performance of our model is not equal across all products in EOL inventories. In particular, we find that demand forecasts for slower-moving products are especially hard to determine. As EOL inventories can hold large quantities of slower-moving SKUs, our findings motivate an extension to our model.

2.3.3 Multiple Series Extension

As evidenced by the work in Section 2.3.1, the accuracy of the parameter estimates of our model depends on attributes of the demand series. Most notably, demand volume was found to be related to the accuracy of parameter estimates. Given the prevalence of slower-moving products in the EOL inventories, we propose a model extension to increase the accuracy of parameter estimates. Based on the vast nature of spare part inventories, we assume there exist common rates of decline in demand amongst SKUs, notwithstanding a difference in volume sizes. Sales history of SKUs with homogeneous rates of decline can be used to estimate the parameters for a single series.

Consider m different SKU demand series each with record length n, where for series i, where $i = 1, ..., m$, the initial mean demand is denoted by $\lambda_{0,i}$ and ρ is the shared rate of decline across a group of series, the demand for series i in period $t = 1, ..., n$ is given by

$$Y_{t,i} \sim \text{ Poisson}(\lambda_{t,i}), \text{ where } \lambda_{t,i} = \lambda_{t-1,i}\rho = \lambda_{0,i}\rho^t.$$

Estimates from a sample size n ($\hat{\rho}_n$ and $\hat{\lambda}_{0,i,n}$) of the parameters (ρ and $\lambda_{0,i}$) are obtained by applying maximum likelihood estimation on historical demand data. For independent demand observations, $y_{1,1}, ..., y_{n,1}, ..., y_{1,m}, ..., y_{n,m}$, belonging to the m series i= $1, ..., m$ we have the likelihood function

$$\mathcal{L}(\lambda_{0,1}, ..., \lambda_{0,m}, \rho; y_{1,1}, ..., y_{n_1,1}, ..., y_{1,m}, ..., y_{n,m}) = \prod_{i=1}^{m}\prod_{t=1}^{n} \frac{(\lambda_{0,i}\rho^t)^{y_{t,i}} e^{-(\lambda_{0,i}\rho^t)}}{y_{t,i}!}$$

and log-likelihood function

$$\ell(\lambda_{0,1}, ..., \lambda_{0,m}, \rho; y_{1,1}, ..., y_{n_1,1}, ..., y_{1,m}, ..., y_{n,m}) \\ = \sum_{i=1}^{m}\sum_{t=1}^{n} \left[y_{t,i}\ln(\lambda_{0,i}\rho^t) - \lambda_{0,i}\rho^t - \ln(y_{t,i}!)\right]. \quad (2.3.6)$$

The likelihood equations are

$$\frac{1}{\hat{\rho}_n}\sum_{i=1}^{m}\sum_{t=1}^{n} t y_{t,i} = \frac{1-\hat{\rho}_n}{\hat{\rho}_n(1-\hat{\rho}_n^n)}\left(\frac{1-\hat{\rho}_n^n}{(1-\hat{\rho}_n)^2} - \frac{n\hat{\rho}_n^n}{1-\hat{\rho}_n}\right)\sum_{i=1}^{m}\sum_{t=1}^{n} y_{t,i}$$

and

$$\hat{\lambda}_{0,i} = \frac{(1-\hat{\rho}_n)\sum_{t=1}^{n} y_{t,i}}{\hat{\rho}_n(1-\hat{\rho}_n^n)}$$

for $i = 1, ..., m$.

The closed form approximation (when n is large) for the maximum likelihood estimator for the common rate of decline is

$$\hat{\rho}_n^A = 1 - \frac{\sum_{i=1}^{m}\sum_{t=1}^{n} y_{t,i}}{\sum_{i=1}^{m}\sum_{t=1}^{n} t y_{t,i}}$$

and correpsondingly the approximate maximum likelihood estimator for $\hat{\lambda}_{0,i,n}$ is

$$\hat{\lambda}^A_{0,i,n} = \frac{(1 - \hat{\rho}^A_n) \sum_{t=1}^{n} y_{t,i}}{\hat{\rho}^A_n}.$$

When $n \to \infty$, $\hat{\rho}_n$ and $\hat{\rho}^A_n$ tend to the same limit, denoted by $\hat{\rho}$. Similarly, $\hat{\lambda}_{0,i,n}$ and $\hat{\lambda}^A_{0,i,n}$ also tend to the same limit, denoted by $\hat{\lambda}_{0,i}$. The limit for the common rate of decline is

$$\hat{\rho} = 1 - \frac{\sum_{i=1}^{m} \sum_{t=1}^{\infty} y_{t,i}}{\sum_{i=1}^{m} \sum_{t=1}^{\infty} t y_{t,i}}$$

and the limit for the estimator for the initial mean demand size is

$$\hat{\lambda}_{0,i} = \frac{(1 - \hat{\rho}) \sum_{t=1}^{\infty} y_{t,i}}{\hat{\rho}}$$

for $i = 1, ..., m$.

To study the behaviour of the limiting estimator of the common rate of decline, $\hat{\rho}$, we consider

$$\hat{\rho} = 1 - \frac{\sum_{i=1}^{m} \sum_{t=1}^{\infty} Y_{t,i}}{\sum_{i=1}^{m} \sum_{t=1}^{\infty} t Y_{t,i}}$$

where $Y_{t,i} \sim$ Poisson($\lambda_{t,i}$) for $i = 1, ..., m$ and all $t \in \mathbb{N}$. We employ approximations for the mean and variance of $\hat{\rho}$ using results on the behaviour of a ratio of random variables using Taylor series expansion. The second order approximation is a function of the combined initial mean demand sizes. To consider the effect of using identically declining series with both similar and differing volume sizes, we obtain a third-order approximation.

Proposition 3: For independently distributed Poisson demand with a geometrically declining mean, the bias and variance of the limiting maximum likelihood estimate for the common rate of decline across series can be obtained. The third-order Taylor series expansion of the bias and variance of a ratio gives

$$\text{Bias}(\hat{\rho}) = \varphi_2\left(\rho, \sum_{i=1}^{m} \lambda_{0,i}^2, \frac{1}{(\sum_{i=1}^{m} \lambda_{0,i})^3}\right) - \varphi_3\left(\rho, \frac{1}{\sum_{i=1}^{m} \lambda_{0,i}}\right),$$

$$\text{Var}(\hat{\rho}) = \varphi_4\Bigg(\rho, \sum_{i=1}^{m} \lambda_{0,i}^2, \sum_{i=1}^{m} \lambda_{0,i}^3, \sum_{i=1}^{m} \lambda_{0,i}^4, \sum_{i=1}^{m} \sum_{\substack{l=1 \\ l\neq i}}^{m} \lambda_{0,i}\lambda_{0,l}, \sum_{i=1}^{m} \sum_{\substack{l=1 \\ l\neq i}}^{m} \lambda_{0,i}^2\lambda_{0,l}, \sum_{i=1}^{m} \sum_{\substack{l=1 \\ l\neq i}}^{m} \lambda_{0,i}^3\lambda_{0,l}$$
$$\sum_{i=1}^{m} \sum_{\substack{l=1 \\ l\neq i}}^{m} \sum_{\substack{r=1 \\ r\neq i,l}}^{m} \lambda_{0,i}\lambda_{0,l}\lambda_{0,r}, \frac{1}{\sum_{i=1}^{m} \lambda_{0,i}}\Bigg),$$

with the functional forms of φ_2, φ_3 and φ_4 as given in Appendix 2.E.

Proposition 3 leads to the following corollary.

Corollary 3: For multiple series with a common decline rate ρ and initial mean demand sizes $\lambda_{0,i}$, for $i = 1, .., m$, the $\text{Bias}(\hat{\rho}) \to 0$ as $\sum_{i=1}^{m} \lambda_{0,i} \to \infty$.

Corollary 3 shows that the bias of the shared rate of decline reduces as the sum of initial mean demand sizes across series becomes large. The variance expression is more difficult to interpret. We therefore conduct a simulation study to evaluate the variance of the shared decline parameter and explore the effect of volume size in Section 2.5.2.

2.3.4 Forecast Performance Metrics

To assess the performance of forecasts for final purchase decisions we are concerned with total demand across all out-of-sample periods, $n+1, ..., N$. Let $f_{t,i}$ be the forecasted demand for series i in time period t and let F_i be the total forecasted demand across the out-of-sample length, $n+1$ to N, so that $F_i = \sum_{t=n+1}^{N} f_{t,i}$. Similarly, where $y_{t,i}$ is the demand for series i in time period t, let the total demand across the out-of-sample period be D_i where $D_i = \sum_{t=n+1}^{N} y_{t,i}$.

The Percentage Error (PE) and the Absolute Percentage Error (APE) across the out-of-sample period for demand series i are given as

$$PE_i = 100\% \left(\frac{D_i - F_i}{D_i} \right)$$

and

$$APE_i = 100\% \left| \frac{D_i - F_i}{D_i} \right| .$$

When aggregating error across groups of series with inconsistent total demand volume sizes, scale-independent metrics are required. The mean percentage error (MPE) and mean absolute percentage error (MAPE) over M series for which performance is measured are given as

$$MPE = \frac{100\%}{M} \sum_{i=1}^{M} \left(\frac{D_i - F_i}{D_i} \right)$$

and

$$MAPE = \frac{100\%}{M} \sum_{i=1}^{M} \left| \frac{D_i - F_i}{D_i} \right| .$$

2.4 Inventory Management

To calculate order quantities for a final purchase decision we assume a base stock policy with fixed stock level O_i for SKU i. The loss function of a demand distribution, G_L, is defined as the expected number of demand units not satisfied. In the model introduced in Section 2.3 the total demand across multiple periods is the sum of independent

Poisson random variables. As such, the total demand across multiple periods is also a Poisson random variable. The loss function for Poisson demand is

$$G_L(O_i) = \lambda_i g(O_i) - (O_i - \lambda_i)(1 - G(O_i))$$

where λ_i is the mean demand for SKU i and G and g are the cumulative distribution function and the probability mass function of the Poisson distribution, respectively (Zipkin, 2000).

Let $\hat{\lambda}_i$ be the forecasted demand across the out-of-sample period. We find the base stock level O_i for a target fill rate $\gamma\%$ by selecting the smallest integer such that the following inequality holds

$$\hat{\lambda}_i g(O_i) - (O_i - \hat{\lambda}_i)(1 - G(O_i)) - (1 - \gamma/100)\hat{\lambda}_i \leq 0.$$

For approaches to incorporate parameter estimation uncertainty from forecasting into inventory we refer the reader to Prak et al. (2017) and Prak and Teunter (2019).

The achieved fill rate (AFR) for series i is calculated as

$$AFR_i = 100\% \left(1 - \frac{(D_i - O_i)^+}{D_i}\right).$$

To evaluate performance over M SKUs we define the mean achieved fill rate (MAFR) as

$$MAFR = \frac{1}{M}\sum_{i=1}^{M} AFR_i.$$

If X is the set of M achieved fill rates ordered from smallest to largest, the median achieved fill rate (MdAFR) is given as

$$MdAFR = \begin{Bmatrix} X[\frac{M+1}{2}], & \text{if } M \text{ is odd} \\ \frac{X[\frac{M}{2}]+X[\frac{M}{2}+1]}{2} & \text{if } M \text{ is even} \end{Bmatrix}.$$

Furthermore, it is of interest to determine the impact of our decisions on leftover inventory. As total demand volume across the out-of-sample period varies depending on demand size, we select a percentage metric for the excess inventory. The excess inventory percentage (EIP) for series i is

$$EIP_i = 100\% \left(\frac{(O_i - D_i)^+}{D_i} \right).$$

The mean excess inventory percentage (MEIP) over M SKUs is

$$MEIP = \frac{1}{M} \sum_{i=1}^{M} EIP_i.$$

We note that relevant product-level information (e.g., price or cost of stockout) could allow manufacturers to alternatively use weighted metrics to assess performance.

2.5 Simulation Study

In Section 2.3, we obtained theoretical results based on asymptotic properties which show that the accuracy of the estimates for the parameters of our model are affected by the characteristics of the demand series. In our model extension, the accuracy of the shared decline rate parameter was found to be impacted by the demand size of the series used in its estimation. We seek to assess the practical relevance of these insights. By conducting a simulation study across a range of parameter values and estimating these parameters numerically, we analyse how demand volume and the declining pattern of demand affect the performance of our model and investigate the impact of varying lengths of demand history on performance. A simulation study allows us to quantify forecast accuracy and inventory performance in a controlled setting. As such, results may be optimistic, and hence, in Section 2.6, we test our model in an uncontrolled setting that reflects the challenges companies face in practice.

Henceforth, we use the abbreviation 'PD' (representing Poisson Decline) for the implementation of the model introduced in Section 2.3.1 and 'PDm' for the multiple series extension to this model as given in Section 2.3.3, where the parameters of the models are estimated by using the Nelder–Mead optimisation method to numerically maximise the log-likelihood functions (2.3.1) and (2.3.6), respectively.

2.5.1 Experiment Design

We simulate data considering a range of different parameter values to showcase possible demand patterns that might be observed in practice. The values of ρ considered in our experiments are equal to $0.97, 0.98$ and 0.99, in accordance with the long-term demand forecasting task for SKUs in the EOL phase. When ρ is low, demand sizes advance quickly to zero. This is particularly true for smaller values of λ_0. We therefore exclude low values of ρ from this study. We include three values for the initial mean demand size, λ_0, that capture the spectrum of monthly demand volume for SKUs with declining demand, in-keeping with SKUs in the end-of-life phase: 10, 50 and 100.

For each combination of input parameters we draw 10,000 demand series from a Poisson distribution with initial mean λ_0 and decline rate ρ. The length of each demand series is $n + 120$ where n is the length of the in-sample. We note that the time period over which performance is measured could differ in practice across products. Due to the significant length of service periods in many applications, we use an out-of-sample forecast horizon of 120 periods in this investigation. The sum of the forecasted demand found using the estimated parameter values and the true total demand across the out-of-sample is used to calculate the forecast percentage error, absolute percentage error, mean achieved fill rate and mean excess inventory percentage. Order quantities are calculated as specified in Section 2.4.

We select a range of values for n. Considering monthly periods, it is less likely that a company will have more than five years of historical data for which demand for a part has been in decline and therefore we do not include values of n greater than 60. It is possible that a company has a limited length of demand history or that a declining pattern has only been observed for a short period of time. In some cases the history can then be limited to only one or two years. Accordingly, we consider in-sample lengths $n = 12$, $n = 24$ and $n = 60$.

Example demand series included in the simulation are given in Figure 2.5.1. In Figure 2.5.1(a), we observe a demand series with low volume size and slow decline, i.e., the first demand value is drawn from a Poisson distribution with a mean equal to $\rho\lambda_0 = 0.99 \times 10 = 9.9$. Subsequently, the second demand value is drawn from a Poisson distribution with mean equal to $\rho^2\lambda_0 = 0.99^2 \times 10 = 9.8$. Hence, we observe a slowly declining demand pattern over time. In contrast, in Figure 2b, the decline rate is lower, and therefore, demand declines more quickly. Figures 2.5.1(c) and 2.5.1(d) display demand series with high volume size, where $\lambda_0 = 100$. By looking at Figure 2.5.1(d), we note periods of zero demand for high volume series when decline is fast.

We also test the multiple series extension presented in Section 2.3.3. For each of the 10,000 series of random observations of length $n + 120$, drawn for each combination of parameters λ_0 and ρ, we draw additional demand series of random observations of length n with the same λ_0 and ρ. The in-sample of each of the 10,000 generated demand series as well as its additional series are then used to estimate the parameters by maximising the log-likelihood function (2.3.6). As above, forecasts are determined for an out-of-sample length of 120 periods. The total number of series used in the application of

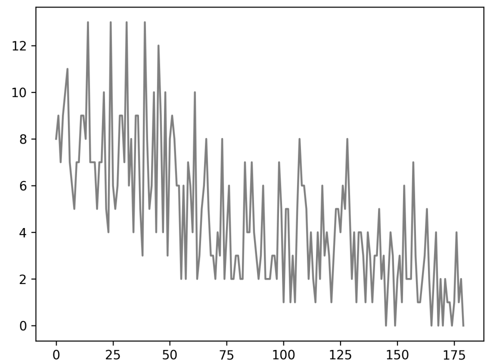

(a) Low Volume, Slow Decline ($\lambda_0 = 10, \rho = 0.99, n = 60$)

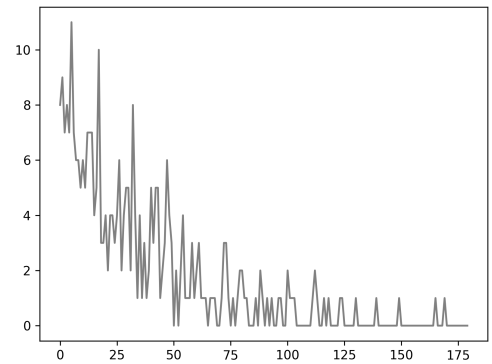

(b) Low Volume, Fast Decline ($\lambda_0 = 10, \rho = 0.97, n = 60$)

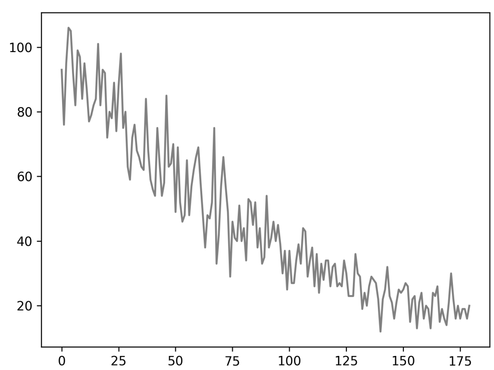

(c) High Volume, Slow Decline ($\lambda_0 = 100, \rho = 0.99, n = 60$)

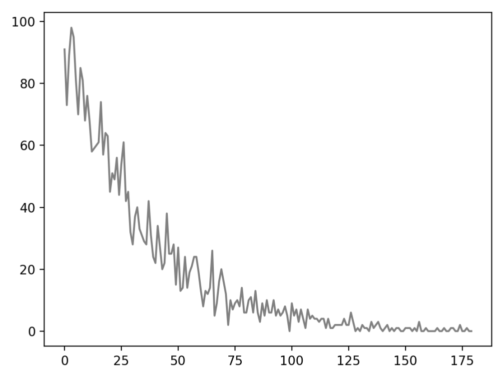

(d) High Volume, Fast Decline ($\lambda_0 = 100, \rho = 0.97, n = 60$)

Figure 2.5.1: Simulated Series Included in Study

PDm is equal to 1 plus the number of additional series. We evaluate the impact of increasing the total initial mean demand size included in the implementation of our multiple series approach. As the initial mean demand sizes in our experimentation are kept consistent this is directly proportional to the number of series used in the multiple series extension to our model.

2.5.2 Results

We give results on the forecast and inventory performance of our model and its extension. We first investigate the performance of using only the demand history from a single series to estimate the parameters. We then evaluate the extension to our model to incorporate information from multiple series with homogeneous decline rates.

Single Series

The best results are obtained when the in-sample length is long. When $n = 60$, the mean percentage error and the mean absolute percentage error are very small. The reduction in mean absolute percentage error is consistent with an increase in the length of the in-sample for all combinations of parameters. We observe that the mean achieved fill rates get closer to the target and the amount of excess inventory reduces as the length of the in-sample grows. An overview of forecast performance for different combinations of initial demand sizes and decline rates are given by Tables 2.5.1 and 2.5.2. Inventory performance is considered in Tables 2.5.3, 2.5.4, 2.5.5. For each combination of parameters, the mean, variance and bias of the parameter estimates over the 10,000 demand series are given in Appendix 2.F.

Table 2.5.1: Mean Percentage Error Using PD for Out-Of-Sample Length of 120 periods (in %)

	$\rho = 0.97$			$\rho = 0.98$			$\rho = 0.99$		
n	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$
12	-67.02	-22.30	-9.72	-43.30	-22.13	-12.81	-8.31	-10.30	-8.57
24	-19.24	-3.56	-1.87	-20.76	-4.59	-2.25	-11.74	-4.77	-2.25
60	-5.59	-1.18	-0.54	-3.92	-0.59	-0.39	-2.92	0.56	-0.28

Demand size impacts the performance of the model. For all lengths of history, we see improvement in the the mean absolute percentage errors as the initial mean demand size increases. The same pattern is clear in the inventory performance whereby the mean achieved fill rate is closer to the target and the mean excess inventory percentage

Table 2.5.2: Mean Absolute Percentage Error Using PD for Out-Of-Sample Length of 120 periods (in %)

	$\rho = 0.97$			$\rho = 0.98$			$\rho = 0.99$		
n	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$
12	111.59	47.35	30.13	91.73	51.93	36.27	61.29	44.16	35.28
24	47.90	18.20	12.90	51.57	20.74	14.14	44.43	22.68	15.69
60	25.32	11.00	7.64	21.48	9.29	6.50	19.01	8.38	5.92

Table 2.5.3: Mean Achieved Fill Rate Using PD for Out-Of-Sample Length of 120 periods and a Target Fill Rate of 95% (in %)

	$\rho = 0.97$			$\rho = 0.98$			$\rho = 0.99$		
n	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$
12	77.10	85.64	87.70	74.92	83.27	86.24	72.54	81.37	84.72
24	85.03	90.49	92.03	83.40	89.69	91.60	82.16	88.81	90.84
60	91.53	93.23	93.90	90.45	93.10	94.04	90.05	93.34	94.16

is smaller when the initial mean demand size is large.

The mean percentage errors are negative in all instances, indicating that the model is over-forecasting. The bias reduces as the in-sample increases. We observe that the rate of decline impacts the bias of the forecasts for longer in-sample lengths. When $n = 60$, the mean percentage errors are closer to zero resulting in a reduced mean excess inventory percentages for series with slower decline patterns than for series with faster decline patterns. We report the mean of the achieved fill rate to demonstrate that the fill rate is closer to the target as the in-sample length and initial mean demand size increases. However, the mean achieved fill rate is sensitive to outlier series which are under-forecasted and achieve low fill rates. Accordingly, we also include the median achieved fill rate in our results. The median achieved fill rates are close to the target

Table 2.5.4: Median Achieved Fill Rate Forecast Using PD for Out-Of-Sample Length of 120 periods and a Target Fill Rate of 95% (in %)

	$\rho = 0.97$			$\rho = 0.98$			$\rho = 0.99$		
n	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$
12	95.97	96.72	95.36	94.41	94.45	94.92	95.40	95.36	95.16
24	97.41	95.81	95.40	96.00	95.42	95.34	96.02	94.83	94.97
60	100.00	96.28	95.43	98.39	95.14	95.08	95.74	95.14	95.00

Table 2.5.5: Mean Excess Inventory Percentage Using PD for Out-Of-Sample Length of 120 periods and a Target Fill Rate of 95% (in %)

	$\rho = 0.97$			$\rho = 0.98$			$\rho = 0.99$		
n	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$
12	82.85	30.66	16.57	61.92	32.80	20.95	30.69	23.44	18.43
24	30.43	8.10	4.80	32.34	9.74	5.55	24.37	10.74	6.30
60	16.13	4.07	2.04	10.89	2.73	1.40	8.39	2.23	1.13

in every case.

Multiple Series

We summarise the effect of our multiple series extension, PDm, where additional series are used to estimate model parameters in order to determine the forecast for a single series. We present two adverse parameter combinations, low volume and fast decline ($\lambda_0 = 10$ and $\rho = 0.97$) and high volume and slow decline ($\lambda_0 = 100$ and $\rho = 0.99$). Noticing that accurate forecasts are more difficult to achieve when historical demand data is limited (as demonstrated in Section 2.5.2), we show results for demand series with a short in-sample of length $n = 12$. Figures 2.5.2, 2.5.3 and 2.5.4 detail the percentage errors, achieved fill rates and excess inventory percentages of the 10,000 forecasts determined for each parameter combination for differing total initial mean demand sizes, $\sum_{i=1}^{m} \lambda_{0,i}$.

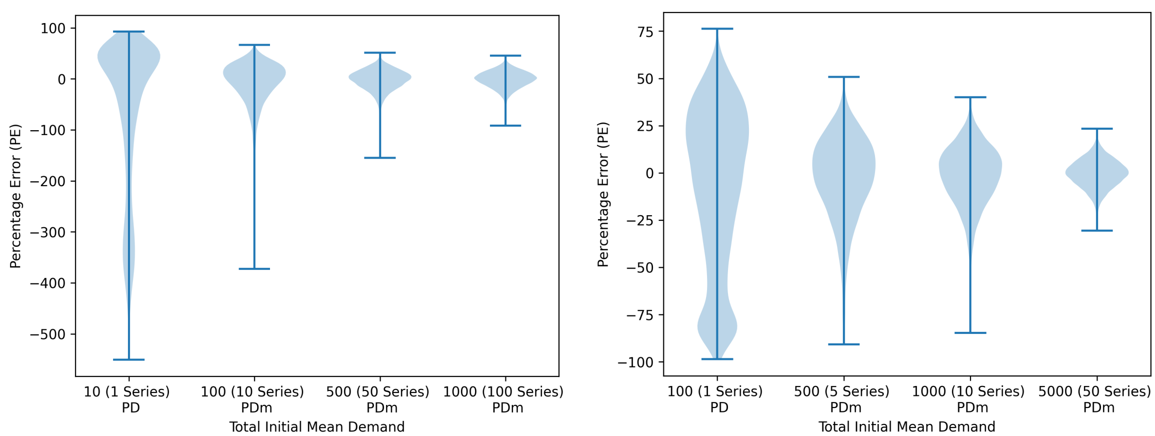

(a) Low Volume, Fast Decline ($\lambda_0 = 10$, $\rho = 0.97, n = 12$)

(b) High Volume, Slow Decline ($\lambda_0 = 100$, $\rho = 0.99, n = 12$)

Figure 2.5.2: Percentage Errors with Increasing Total Initial Mean Demand Size

The multiple series approach improves forecast accuracy and inventory outcomes for SKUs of various volume sizes. The leftmost violin plot in each Sub-Figure within Figures 2.5.2, 2.5.3 and 2.5.4 demonstrates the outcome when using the single series approach, PD. We notice an improvement across all metrics when using PDm over PD.

The larger the total initial mean demand size the more effective the approach is in

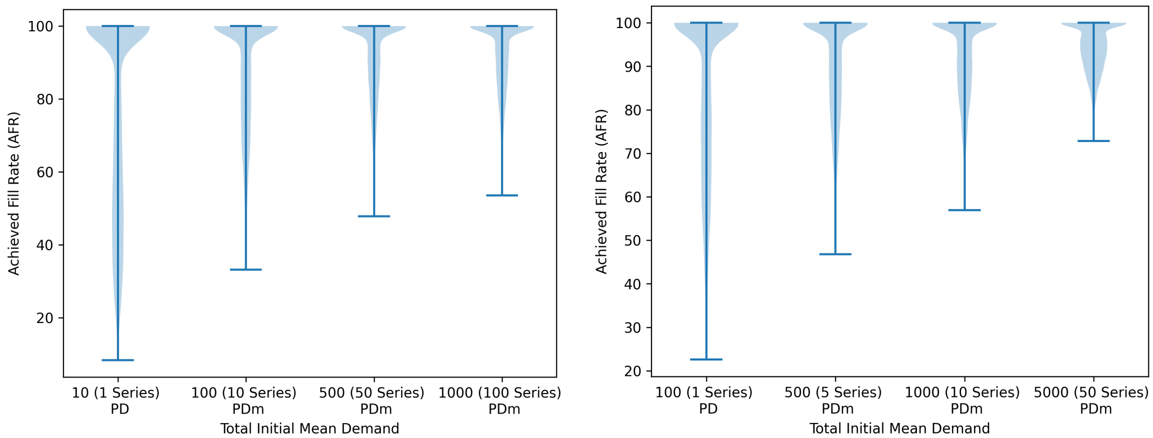

(a) Low Volume, Fast Decline ($\lambda_0 = 10$, $\rho = 0.97, n = 12$)

(b) High Volume, Slow Decline ($\lambda_0 = 100$, $\rho = 0.99, n = 12$)

Figure 2.5.3: Achieved Fill Rates with Increasing Total Initial Mean Demand Size (Target Fill Rate 95%)

improving performance. This insight is in accordance with the analysis in Section 2.3.3. We observe a reduction in variance of the percentage errors, achieved fill rates and the excess inventory percentages as the total initial mean demand size (and the number of series) used to implement PDm increases. These patterns are consistent for low volume and high volume series and series with both slow and fast decline.

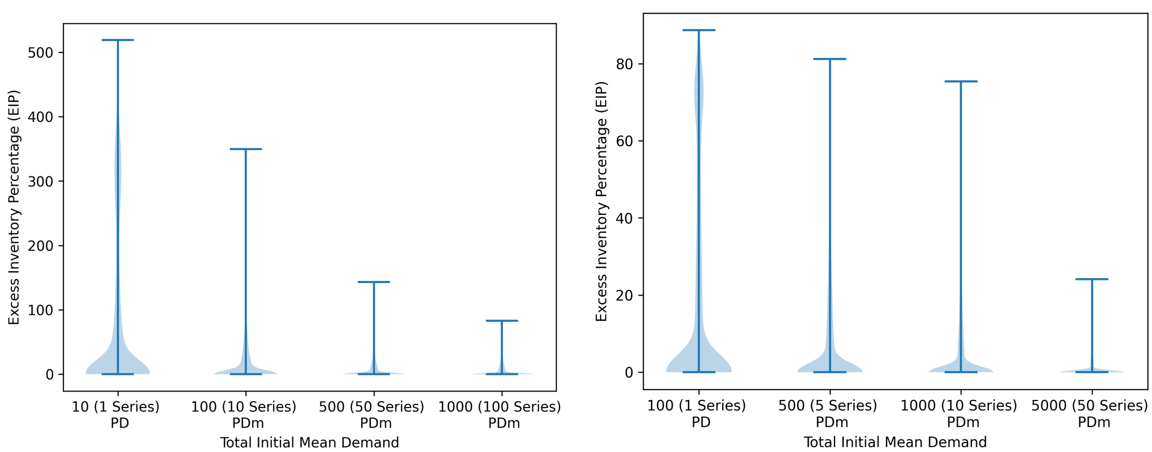

(a) Low Volume, Fast Decline ($\lambda_0 = 10, \rho = 0.97, n = 12$)

(b) High Volume, Slow Decline ($\lambda_0 = 100, \rho = 0.99, n = 12$)

Figure 2.5.4: Excess Inventory with Increasing Total Initial Mean Demand Size (Target Fill Rate 95%)

2.6 Real Data

We conduct a study to show the applicability of our model and multiple series extension using data obtained from an automotive manufacturer.

2.6.1 Overview of Company Data and Benchmark Method

The dataset consists of 9 years of monthly demand data for 400 spare parts over years 2014-2022. The SKUs included in the dataset are in the EOL phase of the life cycle where a sustained decline in demand can be observed. Descriptive statistics of the dataset are given in Table 2.6.1.

Table 2.6.1: Descriptive Statistics of Dataset Consisting of 400 Automotive SKUs

	Mean	Standard Deviation	Min	25th Percentile	Median	75th Percentile	Max
Total Demand	1867.12	3030.40	60	331.00	820.00	1996.25	19556
Number of Periods with Zero Demand	12.81	16.51	0	0.00	4.50	21.00	74

To enable comparison, we include an existing approach in our study as a benchmark. Many existing forecasting approaches are not applicable to our problem setting. Established methods by Holt (1957) and Wright (1986) consider trend in demand and the latter method was adapted for intermittent demand by Altay et al. (2008). However, these approaches were not intended to support long-term decisions and produce one-step ahead forecasts. Additionally, these methods assume a linear trend in demand, contrary to the non-linear declining pattern observed in EOL SKUs, and are therefore unsuitable choices for comparison. More recent forecasting methods suitable for spare parts were proposed for instances were an SKU suddenly becomes obsolete (Teunter et al., 2011; Prestwich et al., 2014, 2021). In these methods, decline in demand is incorporated when a period of zero demand occurs. These methods are designed for short-term forecasting opposed to long-term forecasting where a trend in observed values is expected and produce a flat forecast over the horizon. The forecast procedure outlined by Moore (1971) produces long-term forecasts designed for demand that has a declining long-term forecasting in the EOL phase and considers non-linear patterns of decline. However, this method transforms demand history to a logarithmic scale unsuitable for the SKUs in our study, given that some monthly periods have zero demand.

The approach by Fortuin (1980) is designed for final order decision-making. To the best of our knowledge, a more recent method designed for long-term decision-making concerning SKUs that can have periods of zero demand does not exist. Therefore, we include the approach by Fortuin (1980) in our study to facilitate comparison with our methodology. Fortuin (1980) propose that demand for long-term service periods follows a non-stationary and uncorrelated Gaussian process. They define two parameters, a demand level and a regression factor. Where C, $C > 1$, is the demand level and g, $0 < g < 1$, the regression factor, the mean demand value in period t, μ_t, of the Gaussian process decreases according to

$$\mu_t = Cg^{t-1}.$$

The parameter C is assumed to be known and given by the demand in the period at the start of the EOL period. The regression factor, g, is estimated at 0.7 for the yearly demand of consumer electronic spare parts. We use non-linear least squares to numerically estimate g for the SKUs in our study using demand observations across the in-sample length $1, ..., n$.

In our experimentation, forecasts are evaluated for the total demand over the out-of-sample period. Forecasts are constructed using monthly data over in-sample lengths of 12 and 24 months (1 and 2 years) and evaluated over out-of-sample periods 84 and 96 months (7 and 8 years), respectively. For the implementation of PD, the model introduced in Section 2.3.1, and PDm, our model extension, as given in Section 2.3.3, we estimate the parameters by using the Nelder–Mead optimisation method to numerically maximise the log-likelihood functions (2.3.1) and (2.3.6), respectively. We evaluate performance using the metrics outlined in Sections 2.3.4 and 2.4. We use a base stock policy and a target fill rate of 95% for all approaches. Order quantities are calculated using the loss function for Poisson demand as described in Section 2.4.

In the results presented in Sections 2.6.2 and 2.6.3 we implement PDm using the in-sample data from all SKUs in the dataset to estimate the parameters of the model. In Section 2.6.4 we investigate the effect of using different numbers of SKUs in our multiple series approach.

2.6.2 Overall Performance and In-sample Lengths

Our methodology is effective for forecasting the demand of SKUs in the end-of-life phase. Figure 2.6.1 shows the distribution in performance for all approaches when the in-sample length is 24. The forecast errors, achieved fill rates and excess inventory percentages are more concentrated around their targets for PD than Fortuin's method and are especially good for PDm. When the in-sample length is 24, PDm achieves a MAPE that is 45% lower than Fortuin's method and 35% lower than PD. The improvement in forecast accuracy of our multiple series approach is consistent when the in-sample length is especially short (n=12), in which case the MAPE for PDm is a 65% and 66% less than Fortuin's method and PD, respectively. Table 2.6.2 reports forecast and inventory performance metrics for all approaches. Example forecasts from the study are given in Figure 2.6.2.

The increase in forecast accuracy of PDm leads to better inventory performance. Compared to PD and Fortuin's method, PDm achieves a mean fill rate closer to the target whilst incurring less leftover inventory. When the in-sample length is 24, PDm obtains a 10% and 24% increase in mean achieved fill rate for PD and Fortuin's method, respectively. In line with the results of our controlled study, the achieved fill rates for

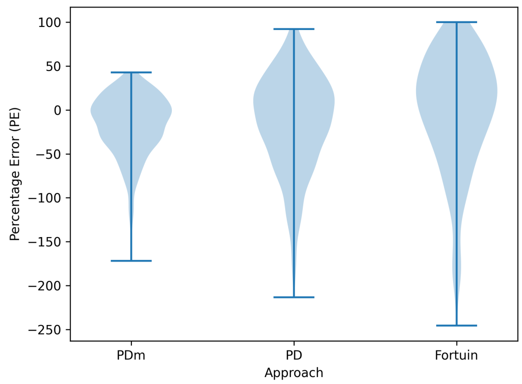

(a) Percentage Error

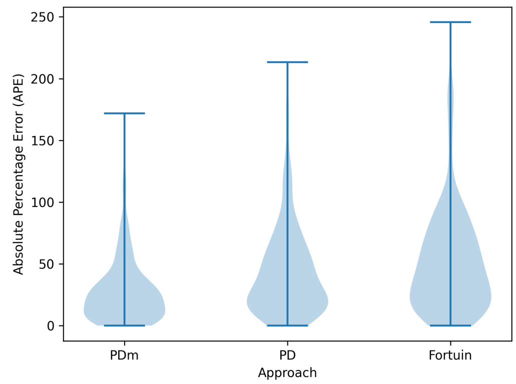

(b) Absolute Percentage Error

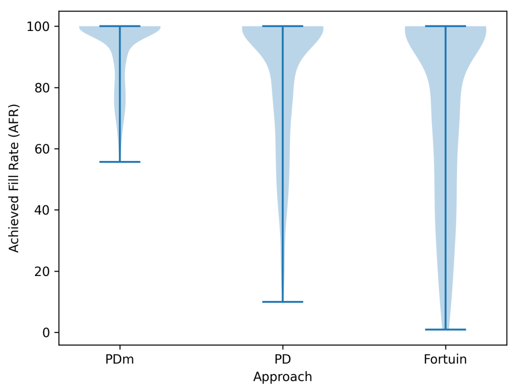

(c) Achieved Fill Rate

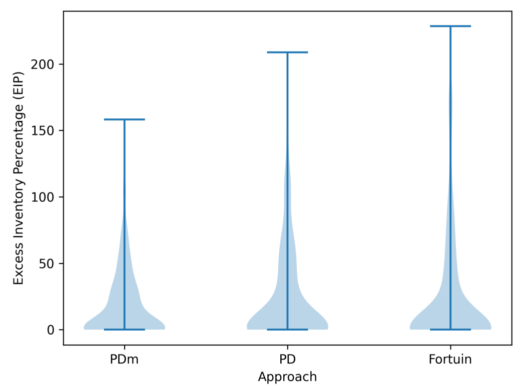

(d) Excess Inventory Percentage

Figure 2.6.1: Forecast and Inventory Performance for 400 Automotive SKUs with $n = 24$ and a Target Fill Rate of 95% (in %)

Table 2.6.2: Forecast and Inventory Performance for 400 Automotive SKUs with a Target Fill Rate of 95% (in %)

	MPE	MAPE	MAFR	MdAFR	MEIP
			$n = 12$		
PDm	9.96	22.10	81.72	83.50	4.68
PD	-8.49	64.54	71.12	79.14	32.76
Fortuin	8.96	63.01	63.27	65.65	24.03
			$n = 24$		
PDm	-17.28	28.43	93.06	100.00	19.47
PD	-14.98	43.92	84.49	100.00	25.84
Fortuin	-2.56	51.59	74.55	87.79	23.98

all methods are closer to the target when the in-sample is 24 than when the in-sample length is 12. The impact of our multiple series approach is realised with respect to both forecast accuracy and inventory decision-making. This shows the benefit of utilising the demand history across inventories when determining forecasts for a single product.

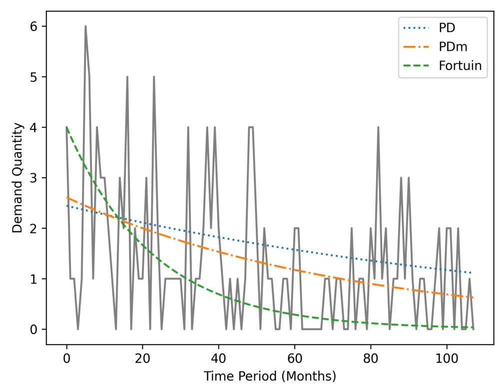

(a) Handle

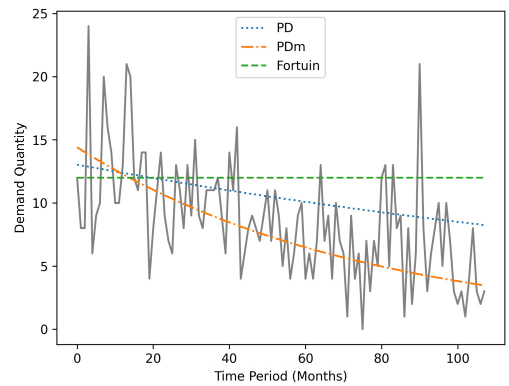

(b) Belt

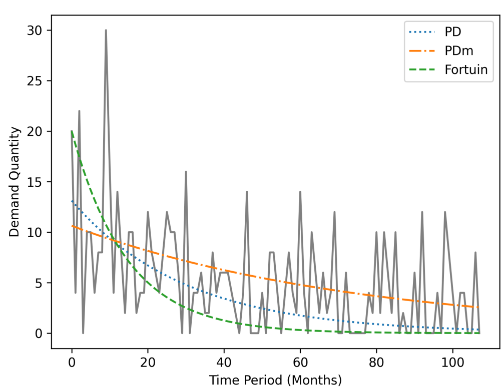

(c) Caliper

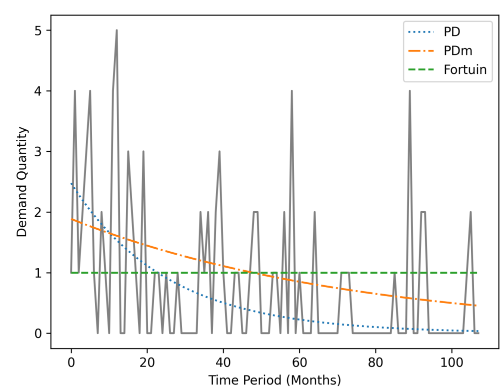

(d) Windscreen

Figure 2.6.2: Examples of Forecasts for SKUs Belonging to an Automotive Manufacturer ($n = 24$)

2.6.3 Slower-moving vs Faster-moving

We analyse the effect of demand size on the performance of our model for real data. As demand for EOL products is non-stationary and declining, we consider the final year of the out-of-sample to analyse forecast performance with respect to demand volume.

An SKU with an annual demand less than or equal to 12 in the final year of the out-of-sample is defined as a slower-moving SKU. Conversely, an SKU with an annual demand of more than 12 in the final year of the out-of-sample is defined as a faster-moving SKU. Using this criteria, 103 SKUs are found to be slower-moving and 297 SKUs are faster-moving.

In keeping with our theoretical findings and the outcomes reported in Section 2.5.2, we find that PD works best for faster-moving SKUs. Table 2.6.3 shows a breakdown of the results from Table 2.6.2 with regard to the demand size of the SKUs. The relationship between demand size and forecast and inventory performance across the SKUs is explored visually in Figure 2.6.3. For PD, the absolute percentage errors are closer to 0 for faster-moving SKUs. The MAPE is lower for faster-moving SKUs than slower-moving SKUs by 30% and 28% for in-sample lengths 12 and 24, respectively. Faster-moving SKUs also achieve fill rates closer to the target. The mean achieved fill rate is 5% closer to the target for faster-moving SKUs than slower-moving SKUs for in-sample length 12 and 2% closer to the target when the in-sample length is 24.

Our multiple series extension improves forecast and inventory performance for both slower-moving SKUs and faster-moving SKUs. In MAPE we see a 75% decrease for slower-moving SKUs and a 61% decrease for faster-moving SKUs when the in-sample length is short (n=12). When the in-sample length is 24 we observe a 32% decrease for slower-moving SKUs and 37% decrease for faster-moving SKUs in MAPE. In Figure 2.6.3c we observe that PDm is effective in improving the achieved fill rates, where the largest benefit is seen with respect to slower-moving SKUs.

Table 2.6.3: Forecast and Inventory Performance for 400 Automotive SKUs Separated into Slow and Fast Moving with a Target Fill Rate of 95% (in %)

	$n = 12$				$n = 24$			
	MPE	MAPE	MAFR	MEIP	MPE	MAPE	MAFR	MEIP
	103 slow-moving SKUs				103 slow-moving SKUs			
PDm	-3.52	20.87	90.44	10.19	-33.91	37.51	97.96	32.29
PD	-19.62	83.16	68.46	47.34	-20.93	55.25	82.93	34.90
Fortuin	7.94	73.79	59.28	30.28	1.98	61.58	68.68	27.30
	297 fast-moving SKUs				297 fast-moving SKUs			
PDm	14.63	22.52	78.69	2.77	-11.51	25.28	91.36	15.02
PD	-4.63	58.08	72.05	27.70	-12.92	40.00	85.03	22.69
Fortuin	9.31	59.28	64.65	21.87	-4.13	48.13	76.58	22.82

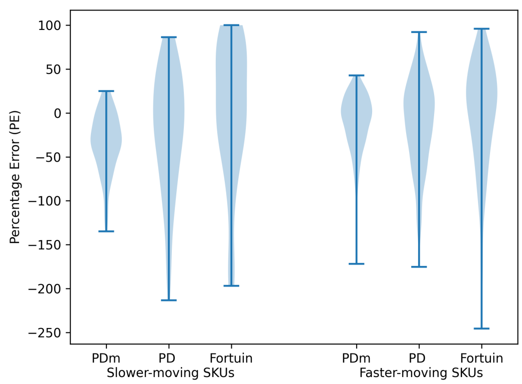

(a) Percentage Error

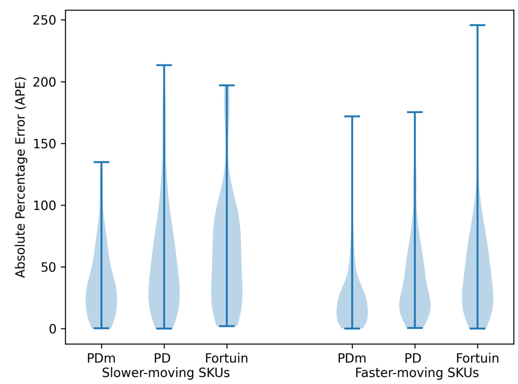

(b) Absolute Percentage Error

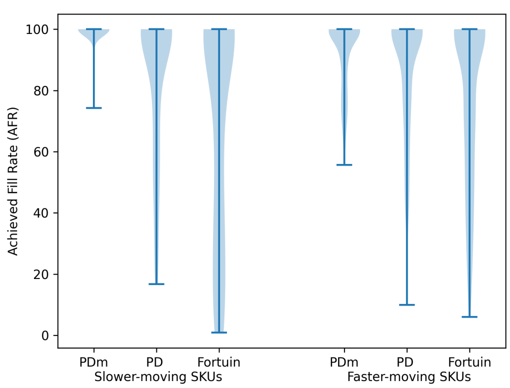

(c) Achieved Fill Rate

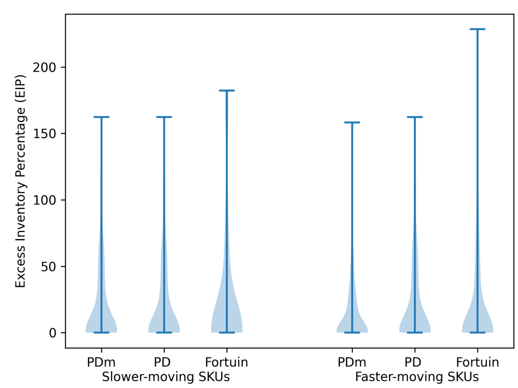

(d) Excess Inventory Percentage

Figure 2.6.3: Forecast and Inventory Performance for (103) Slow-moving SKUs and (297) Fast-moving SKUs In-Sample Length $n = 24$ and a Target Fill Rate of 95%

2.6.4 Number of SKUs in Multiple Series Extension

We use random sampling to enhance the testing of our multiple series approach. For each of the 400 SKUs we randomly select a sample of additional SKUs from the dataset and implement PDm using the SKU's own in-sample demand and the in-sample demand of the additional SKUs from the sample (not the in-sample demand of all SKUs in the dataset) by numerically maximising (2.3.6). We record the percentage error, absolute percentage error, achieved fill rate and excess inventory percentage of the total demand forecast for each SKU. For each of the 400 SKUs, we consider sample sizes of 9, 49, 99 and 199 of randomly selected additional SKUs. The total number of SKUs used to estimate the parameters of the model (m) is then 10, 50, 100 and 200. An evaluation of performance over the 400 SKUs is given in Table 2.6.4.

Table 2.6.4: Forecast and Inventory Performance of Multiple Series Extension (PDm) over 400 Automotive SKUs with a Target Fill Rate of 95% Using Randomly Selected SKUs (in %)

	$n = 12$				$n = 24$			
m	MPE	MAPE	MAFR	MEIP	MPE	MAPE	MAFR	MEIP
10	4.92	37.83	76.90	14.27	-16.08	31.57	90.94	20.44
50	8.70	25.83	80.57	7.02	-16.53	28.99	92.46	19.35
100	9.92	23.36	81.07	5.36	-17.06	28.86	92.73	19.58
200	9.73	22.57	81.55	5.07	-17.13	28.56	92.89	19.50

The number of SKUs considered in the multiple series approach affects forecast accuracy. The MAPE reduces as the number of SKUs in each group increases. Increasing the total number of SKUs used to estimate the model parameters from 10 to 200 incurs a reduction in MAPE of 40% and 10% for in-sample lengths 12 and 24, respectively. The mean achieved fill rate also increases consistently as more SKUs are used in the implementation of PDm. These results support the theoretical insights derived in Section 2.3.3 and highlight the advantage of our multiple series approach.

2.7 Conclusion

We have presented a long-term forecasting model for demand that has a declining trend and that can become intermittent, in line with aftermarket SKUs in the end-of-life phase of the life cycle. We model demand as independent Poisson random variables with a mean that declines geometrically over time. We extend our model to a multiple series scenario, based on the assumption of homogenous decline rates where the demand history of other SKUs is exploited in order to improve estimates and forecasts for a single series.

The parameters of our model are determined via maximum likelihood estimation. Investigation into the bias and variance of the estimates of our model parameters was performed using approximate expressions for the bias and variance of a ratio. We found that demand size was a key factor in the bias and variance of our estimates, where series with larger initial mean demand sizes produced less biased and variable estimates. Our analysis of the bias and variance of approximate estimates of parameters belonging to the extension of our model highlighted the advantage of the sum of initial mean demand across series being large. Numerical results show that estimates and forecasts obtained using the multiple series extension to our model improve upon those made using our single series approach.

We applied the model to a dataset belonging to an automotive manufacturer and found that our approaches work well with respect to forecast accuracy and inventory implications. The single series model performs better for faster-moving than slower-moving SKUs. The multiple series extension to the model is effective in improving forecasts and far outperforms the benchmark, reflected in a lower mean absolute percentage error, a higher mean achieved fill rate and reduced excess inventory.

Our approach can be used to aid other long-term managerial decisions with important environmental and financial implications. Applying the multiple series extension yields more accurate forecasts that better reflect the declining nature of demand, which is particularly relevant for the manufacturer when it comes to long-term planning. Examples at the manufacturer include: calculating the required storage space at any time period over the next ten years to determine whether the existing warehouse space is sufficient or whether a new warehouse is required; reducing CO2 emissions by having fewer parts manufactured and transported to the respective distribution centres; obtaining more accurate long-term demand forecasts that demand planners can use for their financial projections. The model is also applicable to other industries, e.g., retailing, where irregular structures with decline and periods of zero demand can be observed in practice (Sarlo et al., 2023).

In inventories with different decline patterns time-series clustering approaches could be adopted prior to using the multiple series approach. Measures which recognise temporal distortion when finding similarities between series, such as Dynamic Time Warping (Bellman and Kalaba, 1959; Kruskal, 1983; Senin, 2008), could be useful to identify series with similar overall declining patterns when there is variation in demand observations. Alternatively, product information could be useful for grouping parts when they exhibit different declining patterns. Product-specific details held on inventory systems commonly include characteristics such as inventory turnover, the time since a product was introduced to the market and the cost code. The selection of useful grouping attributes is dependent on the industry-or-company-specific context of the application and would require careful consideration by the decision-makers involved.

The focus of this work is on the EOL phase of the life cycle where demand for a part is already in decline. However, it is possible that an LTB decision is required before a declining pattern has been observed. Decline in demand is also assumed to be occurring at a constant rate although in real-life instances the rate of decline may change over time. This work uses the commonly assumed Poisson representation of spare part demand but alternative distributions could be beneficial in the formation of models designed for this problem. For instance, compound Poisson distributions could be utilised to incorporate the variance of non-zero demand sizes. Furthermore, exploring dependence between demand observations could yield further interesting insights, in particular, for practical applications. Another potential area for future research is integrating data from different sources and databases so that causal forecasting methods could be explored. For the example of an automotive manufacturer, this could include bill-of-material information, the number and timing of vehicles sold and the average life spans of different car models. We highlight these aspects as a worthy next steps for research in this area.

2.A Derivation of Maximum Likelihood Equations

For n independent demand observations, $y_1, ..., y_n$, the likelihood function is

$$\mathcal{L}(\lambda_0, \rho; y_1, ..., y_n) = \prod_{t=1}^{n} \frac{(\lambda_0 \rho^t)^{y_t} e^{-(\lambda_0 \rho^t)}}{y_t!}$$

and the log-likelihood function is

$$\ell(\lambda_0, \rho; y_1, ..., y_n) = \sum_{t=1}^{n} \left[y_t \ln(\lambda_0 \rho^t) - \lambda_0 \rho^t - \ln(y_t!) \right].$$

The partial derivative of the log-likelihood function with respect to λ_0 is

$$\frac{\partial \ell}{\partial \lambda_0} = \sum_{t=1}^{n} \left[\frac{y_t}{\lambda_0} - \rho^t \right] \tag{2.A.1}$$

and the partial derivative of the log-likelihood function with respect to ρ is

$$\frac{\partial \ell}{\partial \rho} = \frac{1}{\rho} \sum_{t=1}^{n} [t y_t] - \lambda_0 \sum_{t=1}^{n} t \rho^{t-1}. \tag{2.A.2}$$

By setting (2.A.1) and (2.A.2) equal to zero and rearranging, we have

$$\hat{\lambda}_{0,n} = \frac{\sum_{t=1}^{n} y_t}{\sum_{t=1}^{n} \hat{\rho}_n^t} \tag{2.A.3}$$

and

$$\sum_{t=1}^{n} t\hat{\rho}_n^t = \frac{1}{\hat{\lambda}_{0,n}} \sum_{t=1}^{n} t y_t. \tag{2.A.4}$$

As $\sum_{t=1}^{n} \rho^t = \frac{\rho(1-\rho^n)}{1-\rho}$, we can write (2.A.3) as

$$\hat{\lambda}_{0,n} = \frac{(1-\hat{\rho_n}) \sum_{t=1}^{n} y_t}{\hat{\rho_n}(1-\hat{\rho_n}^n)}$$

and (2.A.4) as

$$\sum_{t=1}^{n} t\hat{\rho}_n^t = \frac{\hat{\rho}_n(1-\hat{\rho}_n^n) \sum_{t=1}^{n} t y_t}{(1-\hat{\rho}_n) \sum_{t=1}^{n} y_t}. \tag{2.A.5}$$

Noticing that

$$\sum_{t=1}^{n} t\rho^t = \frac{\rho\,(1-\rho^n)}{(1-\rho)^2} - \frac{n\rho^{n+1}}{(1-\rho)}$$

(2.A.5) becomes

$$\frac{\hat{\rho}_n\,(1-\hat{\rho}_n^n)}{(1-\hat{\rho}_n)^2} - \frac{n\hat{\rho}_n^{n+1}}{(1-\hat{\rho}_n)} = \frac{\hat{\rho}_n(1-\hat{\rho}_n^n) \sum_{t=1}^{n} t y_t}{(1-\hat{\rho}_n) \sum_{t=1}^{n} y_t}.$$

By simplifying and moving all terms to the left-hand side, we obtain

$$\frac{1}{1-\hat{\rho}_n} - \frac{n\hat{\rho}_n^n}{1-\hat{\rho}_n^n} - \frac{\sum_{t=1}^{n} t y_t}{\sum_{t=1}^{n} y_t} = 0.$$

For a sample of size n, the maximum likelihood estimates are therefore given by solving

$$\hat{\lambda}_{0,n} = \frac{(1-\hat{\rho}_n)\sum_{t=1}^{n} y_t}{\hat{\rho}_n(1-\hat{\rho}_n{}^n)}$$

and

$$\frac{1}{1-\hat{\rho}_n} - \frac{n\hat{\rho}_n{}^n}{1-\hat{\rho}_n{}^n} - \frac{\sum_{t=1}^{n} t y_t}{\sum_{t=1}^{n} y_t} = 0.$$

2.B Limit of Fisher Information

For a sample size n, the second partial derivatives of the log-likelihood function ℓ, denoted here by ℓ_n to emphasise its dependence on n, with respect to parameters λ_0 and ρ are given by

$$\frac{\partial^2 \ell_n}{\partial \lambda_0^2} = -\frac{1}{\lambda_0^2}\sum_{t=1}^{n} y_t,$$

$$\frac{\partial^2 \ell_n}{\partial \rho^2} = -\frac{1}{\rho^2}\sum_{t=1}^{n} t y_t - \lambda_0 \sum_{t=2}^{n} t(t-1)\rho^{t-2},$$

and

$$\frac{\partial^2 \ell_n}{\partial \lambda_0 \partial \rho} = -\sum_{t=1}^{n} t\rho^{t-1}.$$

The elements of the Fisher Information matrix are

$$\mathbb{E}\left(-\frac{\partial^2 \ell_n}{\partial \lambda_0^2}\right) = \frac{\rho(1-\rho^n)}{\lambda_0(1-\rho)},$$

$$\begin{aligned}\mathbb{E}\left(-\frac{\partial^2 \ell_n}{\partial \rho^2}\right) &= \frac{\lambda_0}{\rho}\left(\frac{\rho\,(1-\rho^n)}{(1-\rho)^2} + \frac{1-(n+1)\rho^n}{(1-\rho)}\right) \\ &\quad + \lambda_0\left(\frac{-(n+1)n\rho^{n-1}}{(1-\rho)} + \frac{2(1-(n+1))\rho^n}{(1-\rho)^2} + \frac{2\rho(1-\rho^n)}{(1-\rho)^3}\right),\end{aligned}$$

and

$$\mathbb{E}\left(-\frac{\partial^2 \ell_n}{\partial \lambda_0 \partial \rho}\right) = \frac{1-(n+1)\rho^n}{1-\rho} + \frac{\rho(1-\rho^n)}{(1-\rho)^2}.$$

As $n \to \infty$,

$$\mathbb{E}\left(-\frac{\partial^2 \ell_n}{\partial \lambda_0^2}\right) \to \frac{\rho}{\lambda_0(1-\rho)},$$

$$\mathbb{E}\left(-\frac{\partial^2 \ell_n}{\partial \rho^2}\right) \to \frac{\lambda_0}{\rho}\left(\frac{1}{(1-\rho)^2}\right) + \lambda_0\left(\frac{2}{(1-\rho)^3}\right) = \frac{\lambda_0(1+\rho)}{\rho(1-\rho)^3},$$

and

$$\mathbb{E}\left(-\frac{\partial^2 \ell_n}{\partial \lambda_0 \partial \rho}\right) \to \frac{1}{(1-\rho)^2}.$$

We observe that, as $n \to \infty$, the elements of the Fisher Information matrix tend to finite limits.

2.C Proof of Proposition 1

The Taylor series approximation for the expectation of a ratio truncated at the second term is

$$\mathbb{E}\left(\frac{Q}{R}\right) \approx \frac{\mathbb{E}(Q)}{\mathbb{E}(R)}\left(1 - \frac{\mathrm{Cov}(Q,R)}{\mathbb{E}(Q)\mathbb{E}(R)} + \frac{\mathrm{Var}(R)}{\mathbb{E}(R)^2}\right).$$

The expression for the expectation of the limiting estimator for the rate of decline, $\hat{\rho}$, is given by

$$\mathbb{E}(\hat{\rho}) = 1 - \mathbb{E}\left(\frac{\sum_{t=1}^{\infty} Y_t}{\sum_{t=1}^{\infty} tY_t}\right).$$

We calculate the necessary components:

$$\frac{\mathbb{E}\left(\sum_{t=1}^{\infty} Y_t\right)}{\mathbb{E}\left(\sum_{t=1}^{\infty} tY_t\right)} = 1-\rho,$$

$$\frac{\text{Cov}\left(\sum_{t=1}^{\infty} Y_t, \sum_{t=1}^{\infty} tY_t\right)}{\mathbb{E}\left(\sum_{t=1}^{\infty} Y_t\right)\mathbb{E}\left(\sum_{t=1}^{\infty} tY_t\right)} = \frac{1-\rho}{\lambda_0\rho},$$

$$\frac{\text{Var}(\sum_{t=1}^{\infty} tY_t)}{\mathbb{E}\left(\sum_{t=1}^{\infty} tY_t\right)^2} = \frac{(\rho+\rho^2)(1-\rho)}{\lambda_0\rho^2}.$$

The approximate expectation of $\hat{\rho}$ is

$$\mathbb{E}(\hat{\rho}) \approx 1 - \left(1-\rho+\frac{(1-\rho)^2}{\lambda_0}\right).$$

It is easy to see that the approximate expression for the bias of $\hat{\rho}$ can be written as

$$Bias(\hat{\rho}) \approx -\frac{(1-\rho)^2}{\lambda_0}.$$

The Taylor series approximation for the variance of a ratio truncated at the second term is given by

$$\text{Var}\left(\frac{Q}{R}\right) \approx \frac{\mathbb{E}(Q)^2}{\mathbb{E}(R)^2}\left(\frac{\text{Var}(Q)}{\mathbb{E}(Q)^2} - \frac{2\text{Cov}(Q,R)}{\mathbb{E}(Q)\mathbb{E}(R)} + \frac{\text{Var}(R)}{\mathbb{E}(R)^2}\right).$$

The variance of $\hat{\rho}$ is given as

$$\text{Var}(\hat{\rho}) = \text{Var}\left(1-\frac{\sum_{t=1}^{\infty} Y_t}{\sum_{t=1}^{\infty} tY_t}\right) = \text{Var}\left(\frac{\sum_{t=1}^{\infty} Y_t}{\sum_{t=1}^{\infty} tY_t}\right).$$

We write

$$\frac{\text{Var}(\sum_{t=1}^{\infty} Y_t)}{\mathbb{E}(\sum_{t=1}^{\infty} Y_t)^2} = \frac{1-\rho}{\lambda_0\rho},$$

$$\frac{\mathbb{E}(\sum_{t=1}^{\infty} Y_t)^2}{\mathbb{E}(\sum_{t=1}^{\infty} tY_t)^2} = (1-\rho)^2,$$

$$-\frac{2\mathrm{Cov}(\sum_{t=1}^{\infty} Y_t, \sum_{t=1}^{\infty} tY_t)}{\mathbb{E}(\sum_{t=1}^{\infty} Y_t)\mathbb{E}(\sum_{t=1}^{\infty} tY_t)} = -\frac{2(1-\rho)}{\lambda_0\rho},$$

$$\frac{\mathrm{Var}(\sum_{t=1}^{\infty} tY_t)}{\mathbb{E}(\sum_{t=1}^{\infty} tY_t)^2} = \frac{1-\rho^2}{\lambda_0\rho}.$$

The approximate expression for the variance of $\hat{\rho}$ can therefore be written as

$$\mathrm{Var}(\hat{\rho}) \approx \frac{(1-\rho)^3}{\lambda_0}.$$

2.D Proof of Proposition 2

The Taylor series approximation for the expectation of a ratio truncated at the second term is

$$\mathbb{E}\left(\frac{Q}{R}\right) \approx \frac{\mathbb{E}(Q)}{\mathbb{E}(R)}\left(1 - \frac{\mathrm{Cov}(Q,R)}{\mathbb{E}(Q)\mathbb{E}(R)} + \frac{\mathrm{Var}(R)}{\mathbb{E}(R)^2}\right).$$

The expectation of the limiting estimator for the initial mean demand size, $\hat{\lambda}_0$, is given by

$$\mathbb{E}(\hat{\lambda}_0) = \mathbb{E}\left(\frac{(\sum_{t=1}^{\infty} Y_t)^2}{\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t}\right).$$

We obtain

$$\mathbb{E}\left(\sum_{t=1}^{\infty} Y_t\right)^2 = \frac{\lambda_0\rho}{1-\rho} + \frac{\lambda_0^2\rho^2}{(1-\rho)^2},$$

$$\mathbb{E}\left(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right) = \frac{\lambda_0\rho}{(1-\rho)^2} - \frac{\lambda_0\rho}{1-\rho},$$

$$\mathrm{Cov}\left(\left(\sum_{t=1}^{\infty} Y_t\right)^2, \sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right) = \frac{2\lambda_0^2\rho^2}{(1-\rho)^3} - \frac{\lambda_0^2\rho^2}{(1-\rho)^2} - \frac{\lambda_0\rho}{1-\rho},$$

$$\mathrm{Var}\left(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right) = \frac{\lambda_0(\rho+\rho^2)}{(1-\rho)^3} + -\frac{2\lambda_0\rho}{(1-\rho)^2} + \frac{\lambda_0\rho}{1-\rho},$$

$$\mathrm{Var}\left(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right) = \mathrm{Var}\left(\sum_{t=1}^{\infty}(t-1)Y_t\right) = \frac{\lambda_0(\rho+\rho^2)}{(1-\rho)^3} + -\frac{2\lambda_0\rho}{(1-\rho)^2} + \frac{\lambda_0\rho}{1-\rho},$$

$$\frac{\mathbb{E}\left(\sum_{t=1}^{\infty} Y_t\right)^2}{\mathbb{E}\left(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right)} = \lambda_0 + \frac{1-\rho}{\rho},$$

$$\mathbb{E}\left(\sum_{t=1}^{\infty} Y_t\right)^2 \mathbb{E}\left(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right) = \frac{\rho^3\lambda_0^2(1+\lambda_0\rho-\rho)}{(1-\rho)^4},$$

$$\frac{\mathrm{Var}(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t)}{\mathbb{E}(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t)^2} = \frac{1-\rho^2}{\lambda_0\rho^2},$$

$$\frac{\mathrm{Cov}\left((\sum_{t=1}^{\infty} Y_t)^2, \sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right)}{\mathbb{E}\left(\sum_{t=1}^{\infty} Y_t\right)^2 \mathbb{E}\left(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right)} = \frac{(1-2\lambda_0\rho)(1-\rho)^2 + (1-\rho)2\lambda_0\rho - (1-\rho)^3}{\rho^2\lambda_0(1+\lambda_0\rho-\rho)}.$$

The approximate expression for the bias of $\hat{\lambda}_0$ is denoted by

$$\mathrm{Bias}(\hat{\lambda}_0) \approx \frac{(1-\rho)^2}{\rho^3\lambda_0} + \frac{1-\rho}{\rho^2}.$$

The Taylor series approximation for the variance of a ratio truncated at the second term is given by

$$\text{Var}\left(\frac{Q}{R}\right) \approx \frac{\mathbb{E}(Q)^2}{\mathbb{E}(R)^2}\left(\frac{\text{Var}(Q)}{\mathbb{E}(R)^2} - \frac{2\text{Cov}(Q,R)}{\mathbb{E}(Q)\mathbb{E}(R)} + \frac{\text{Var}(R)}{\mathbb{E}(R)^2}\right).$$

The variance of $\hat{\lambda}_0$ is given by

$$Var(\hat{\lambda}_0) = Var\left(\frac{(\sum_{t=1}^{\infty} Y_t)^2}{\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t}\right).$$

We obtain

$$\mathbb{E}\left(\left(\sum_{t=1}^{\infty} Y_t\right)^2\right)^2 = \frac{\lambda_0^2\rho^2(\lambda_0\rho - \rho + 1)^2}{(1-\rho)^4},$$

$$\mathbb{E}\left(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right)^2 = \frac{\lambda_0^2\rho^4}{(1-\rho)^4},$$

$$\begin{aligned}\text{Var}\left(\left(\sum_{t=1}^{\infty} Y_t\right)^2\right) = &\frac{\lambda_0^4\rho^4}{1-\rho^4} + \frac{6\lambda_0^3\rho^3}{1-\rho^3} + \frac{7\lambda_0^2\rho^2}{1-\rho^2} + \frac{\lambda_0\rho}{1-\rho} + \frac{8\lambda_0^4\rho^5}{(1-\rho^4)(1-\rho)} \\ &+ \frac{48\lambda_0^3\rho^4}{(1-\rho^3)(1-\rho)} + \frac{20\lambda_0^2\rho^3}{(1-\rho)(1-\rho^2)} + \frac{12\lambda_0^4\rho^6}{(1-\rho^2)(1-\rho^4)} \\ &+ \frac{36\lambda_0^4\rho^7}{(1-\rho)(1-\rho^2)(1-\rho^4)} + \frac{36\lambda_0^3\rho^6}{(1-\rho)(1-\rho^2)(1-\rho^3)} \\ &+ \frac{96\lambda_0^4\rho^{10}}{(1-\rho)(1-\rho^2)(1-\rho^3)(1-\rho^4)} \\ &- \frac{\lambda_0^2\rho^2(\lambda_0\rho - \rho + 1)^2}{(1-\rho)^4},\end{aligned}$$

$$\frac{\mathbb{E}\left((\sum_{t=1}^{\infty} Y_t)^2\right)^2}{\mathbb{E}(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t)^2} = \frac{(\lambda_0\rho - \rho + 1)^2}{\rho^2},$$

$$\frac{\mathrm{Var}\left(\left(\left(\sum_{t=1}^{\infty} Y_t\right)^2\right)^2\right)}{\mathbb{E}(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t)^2} = \frac{(1-\rho)^4}{\rho^4}\Bigg(\frac{\lambda_0^2\rho^4}{1-\rho^4} + \frac{6\lambda_0\rho^3}{1-\rho^3} + \frac{7\rho^2}{1-\rho^2} + \frac{\rho}{\lambda_0(1-\rho)}$$
$$+ \frac{8\lambda_0^2\rho^5}{(1-\rho^4)(1-\rho)} + \frac{48\lambda_0\rho^4}{(1-\rho^3)(1-\rho)} + \frac{20\rho^3}{(1-\rho)(1-\rho^2)}$$
$$+ \frac{12\lambda_0^2\rho^6}{(1-\rho^2)(1-\rho^4)} + \frac{36\lambda_0^2\rho^7}{(1-\rho)(1-\rho^2)(1-\rho^4)}$$
$$+ \frac{36\lambda_0\rho^6}{(1-\rho)(1-\rho^2)(1-\rho^3)} + \frac{96\lambda_0^2\rho^{10}}{(1-\rho)(1-\rho^2)(1-\rho^3)(1-\rho^4)}$$
$$- \frac{\rho^2(\lambda_0\rho-\rho+1)^2}{(1-\rho)^4}\Bigg), \tag{2.D.1}$$

$$\mathbb{E}\left(\left(\sum_{t=1}^{\infty} Y_t\right)^2\right)\mathbb{E}\left(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right) = \frac{\rho^3\lambda_0^2(1+\lambda_0\rho-\rho)}{(1-\rho)^4},$$

$$\frac{-2\mathrm{Cov}\left(\left(\sum_{t=1}^{\infty} Y_t\right)^2, \sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right)}{\mathbb{E}\left(\left(\sum_{t=1}^{\infty} Y_t\right)^2\right)\mathbb{E}\left(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right)} = \frac{-4\lambda_0\rho(1-\rho) + 2\lambda_0\rho(1-\rho)^2 + (1-\rho)^3}{\rho^2\lambda_0(1+\lambda_0\rho-\rho)},$$

$$\frac{Var(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t)}{\mathbb{E}(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t)^2} = \frac{(\rho+\rho^2)(1-\rho) - 2\rho(1-\rho)^2 + \rho(1-\rho)^3}{\lambda_0\rho^4},$$

$$\frac{-2\mathrm{Cov}((\sum_{t=1}^{\infty} Y_t)^2, \sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t)\mathbb{E}((\sum_{t=1}^{\infty} Y_t)^2)}{\mathbb{E}(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t)^3}$$
$$= \frac{(\lambda_0\rho-\rho+1)^2(-4\lambda_0\rho(1-\rho) + 2\lambda_0\rho(1-\rho)^2 + (1-\rho)^3)}{\rho^4\lambda_0(1+\lambda_0\rho-\rho)}, \tag{2.D.2}$$

$$\frac{\mathbb{E}\left(\left(\sum_{t=1}^{\infty} Y_t\right)^2\right)^2 Var\left(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right)}{\mathbb{E}\left(\sum_{t=1}^{\infty} tY_t - \sum_{t=1}^{\infty} Y_t\right)^4}$$
$$= \frac{(\lambda_0\rho - \rho + 1)^2((\rho + \rho^2)(1-\rho) - 2\rho(1-\rho)^2 + \rho(1-\rho)^3)}{\lambda_0\rho^6}. \quad (2.D.3)$$

The approximate expression for the variance of $\hat{\lambda}_0$ is given as the sum of (2.D.1), (2.D.2) and (2.D.3).

2.E Proof of Proposition 3

The Taylor series approximations for the bias and variance of a ratio truncated at the third term are

$$\mathbb{E}\left(\frac{Q}{R}\right) \approx \frac{\mathbb{E}(Q)}{\mathbb{E}(R)}\left(1 - \frac{\mathrm{Cov}(Q,R)}{\mathbb{E}(Q)\mathbb{E}(R)} + \frac{\mathrm{Var}(R)}{\mathbb{E}(R)^2} + \frac{Cov(Q,R^2)}{\mathbb{E}(Q)\mathbb{E}(R)^2} - \frac{\mu_3(R)}{\mathbb{E}(R)^3}\right)$$

and

$$Var\left(\frac{Q}{R}\right) \approx \frac{\mathbb{E}(Q)}{\mathbb{E}(R)}\left(\frac{Var(Q)}{\mathbb{E}(Q)\mathbb{E}(R)} + \frac{\mathbb{E}(Q)Var(R)}{\mathbb{E}(R)^3} + \frac{-2Cov(Q,R)}{\mathbb{E}(R)^2} + \frac{-2Cov(Q^2,R)}{\mathbb{E}(Q)\mathbb{E}(R)^2}\right.$$
$$\left. + \frac{-2\mu_3(R)\mathbb{E}(Q)}{E(R)^4} + \frac{Cov(Q^2,R^2)}{E(Q)E(R)^3} + \frac{-2Cov(Q,R^3)}{E(R)^4} + \frac{\mu_4(R)}{2\mathbb{E}(R)^2}\right)$$

where μ_3 and μ_4 are the third and fourth central moment, respectively.

The expectation of the limiting estimator for the rate of decline, $\hat{\rho}$, shared across m series is given by

$$\mathbb{E}(\hat{\rho}) = \mathbb{E}\left(1 - \frac{\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i}}{\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i}}\right) = 1 - \mathbb{E}\left(\frac{\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i}}{\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i}}\right).$$

We obtain

$$\frac{\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i})}{\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})} = 1-\rho,$$

$$-\frac{\text{Cov}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i}, \sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})}{(\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i}))^2} = -\frac{(1-\rho)^2}{\rho\sum_{i=1}^{m}\lambda_{0,i}},$$

$$\frac{\text{Var}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i})}{(\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i}))^2} = \frac{(1-\rho)(1-\rho^2)}{\rho\sum_{i=1}^{m}\lambda_{0,i}},$$

$$\begin{aligned}\frac{Cov\left(\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i}, (\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})^2\right)}{\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})^3} &= \frac{-2(1+\rho^2)(1-\rho)^6\sum_{i=1}^{m}\lambda_{0,i}^2}{\rho(1-\rho^2)^3(\sum_{i=1}^{m}\lambda_{0,i})^3} \\ &+ \frac{(1+\rho)(1-\rho)^3}{\rho^2(\sum_{i=1}^{m}\lambda_{0,i})^2} + \frac{4(1-\rho)^2}{\rho(\sum_{i=1}^{m}\lambda_{0,i})},\end{aligned}$$

$$-\frac{\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i})\mu_3(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})}{\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})^4} = -\frac{(1-\rho)^3(\rho^2+4\rho+1)}{\rho^2(\sum_{i=1}^{m}\lambda_{0,i})^2}.$$

The approximate expression for the bias of $\hat{\rho}$ in the multiple series setting is given by

$$\begin{aligned}Bias(\hat{\rho}) \approx &-\frac{(\rho-1)^2}{\sum_{i=1}^{m}\lambda_{0,i}} + \frac{2(1+\rho^2)(1-\rho)^6\sum_{i=1}^{m}\lambda_{0,i}^2}{\rho(1-\rho^2)^3(\sum_{i=1}^{m}\lambda_{0,i})^3} - \frac{(1-\rho)^3(-\rho^2-3\rho)}{\rho^2(\sum_{i=1}^{m}\lambda_{0,i})^2} \\ &- \frac{4(1-\rho)^2}{\rho(\sum_{i=1}^{m}\lambda_{0,i})}.\end{aligned}$$

The approximate expression for the variance of $\hat{\rho}$ common across m series is given as

$$\mathrm{Var}(\hat{\rho}^A) = \mathrm{Var}\left(1 - \frac{\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i}}{\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i}}\right) = \mathrm{Var}\left(\frac{\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i}}{\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i}}\right).$$

We obtain

$$\frac{\mathrm{Var}(\sum_{t=1}^{m}\sum_{t=1}^{\infty} Y_{t,i})}{\mathbb{E}(\sum_{t=1}^{m}\sum_{t=1}^{\infty} Y_{t,i})^2}\frac{\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} y_{t,j})^2}{\mathbb{E}(\sum_{t=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})^2} = \frac{(1-\rho)^3}{\rho\sum_{i=1}^{m}\lambda_{0,i}}, \tag{2.E.1}$$

$$\frac{\mathrm{Var}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})\mathbb{E}(\sum_{i=1}^{m}\sum_{i=1}^{\infty} Y_{t,i})^2}{\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})^4} = \frac{(1-\rho)^3(1+\rho)}{\rho\sum_{i=1}^{m}\lambda_{0,i}}, \tag{2.E.2}$$

$$\frac{-2\mathrm{Cov}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i},\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})}{\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i})\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})}\frac{\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i})^2}{\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})^2} = \frac{-2(1-\rho)^3}{\rho\sum_{i=1}^{m}\lambda_{0,i}}, \tag{2.E.3}$$

$$\frac{-2\mu_3(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})\mathbb{E}(\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i})^2}{E(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})^5} = \frac{-2(\rho^2+4\rho+1)(1-\rho)^4}{\rho^2(\sum_{i=1}^{m}\lambda_{0,i})^2}, \tag{2.E.4}$$

$$\begin{aligned}\frac{-2Cov\left((\sum_{i=1}^{m}\sum_{t=1}^{\infty} Y_{t,i})^2,(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i})\right)}{\mathbb{E}\left(\sum_{i=1}^{m}\sum_{t=1}^{\infty} tY_{t,i}\right)^3} &= -\frac{2}{\rho(\sum_{j=1}^{m}\lambda_{0,j})(1-\rho)^2} - \frac{4}{(1-\rho^2)^2}\\ &\quad - \frac{8}{(1-\rho)^3},\end{aligned} \tag{2.E.5}$$

$$
\begin{aligned}
&\frac{Cov\left(\left(\sum_{i=1}^{m}\sum_{t=1}^{\infty}Y_{t,i}\right)^2,\left(\sum_{i=1}^{m}\sum_{t=1}^{\infty}tY_{t,i}\right)^2\right)}{\mathbb{E}\left(\sum_{i=1}^{m}\sum_{t=1}^{\infty}tY_{t,i}\right)^4} \\
&= \frac{(1-\rho)^8}{\rho^4(\sum_{i=1}^{m}\lambda_{0,i})^4}\left(\sum_{i=1}^{m}\lambda_{0,i}^4\left(8\sum_{t=1}^{\infty}tu\rho^{3t+u}+8\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}t^2\rho^{3t+u}\right)\right. \\
&\quad+\sum_{i=1}^{m}\lambda_{0,i}^3\left(\sum_{t=1}^{\infty}4t^2\rho^{3t}+5\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}tu\rho^{2t+u}+20\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}t^2\rho^{2t+u}\right. \\
&\quad\left.+16\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}\sum_{\substack{v=1\\v\neq u,t}}^{\infty}tu\rho^{t+u+v}\right)+\sum_{i=1}^{m}\lambda_{0,i}^2\left(6\sum_{t=1}^{\infty}t^2\rho^{2t}+4\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}tu\rho^{t+u}\right. \\
&\quad\left.+4\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}t^3\rho^{t+u}\right)+\sum_{i=1}^{m}\lambda_{0,i}\sum_{t=1}^{\infty}t^2\rho^t+\sum_{i=1}^{m}\sum_{\substack{l=1\\l\neq i}}^{m}\lambda_{0,i}^2\lambda_{0,l}\left(40\sum_{t=1}^{\infty}t^2\rho^{3t}\right. \\
&\quad\left.+68\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}tu\rho^{2t+u}+36\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}t^2\rho^{2t+u}+32\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}\sum_{\substack{v=1\\v\neq t,u}}^{\infty}tu\rho^{t+u+v}\right) \\
&\quad+\sum_{i=1}^{m}\sum_{\substack{l=1\\l\neq i}}^{m}\lambda_{0,i}\lambda_{0,l}\left(8\sum_{t=1}^{\infty}t^2\rho^{2t}+12\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}tu\rho^{t+u}+\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}t^2\rho^{t+u}\right) \\
&\quad+\sum_{i=1}^{m}\sum_{\substack{q=1\\l\neq i}}^{m}\lambda_{0,i}^3\lambda_{0,l}\left(8\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}tu\rho^{3t+u}+16\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}t^2\rho^{4t}+8\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}t^2\rho^{3t+u}\right) \\
&\quad+16\sum_{i=1}^{m}\sum_{\substack{l=1\\l\neq i}}^{m}\sum_{\substack{r=1\\r\neq l,i}}^{m}\lambda_{0,i}\lambda_{0,l}\lambda_{0,r}\left(\sum_{t=1}^{\infty}t^2\rho^{3t}+\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}tu\rho^{2t+u}\right. \\
&\quad\left.\left.+\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}t^2\rho^{2t+v}+\sum_{t=1}^{\infty}\sum_{\substack{u=1\\u\neq t}}^{\infty}\sum_{\substack{v=1\\v\neq t,u}}^{\infty}tu\rho^{t+u+v}\right)\right),
\end{aligned}
\tag{2.E.6}
$$

$$
\begin{aligned}
&\frac{-2\mathbb{E}\left(\sum_{i=1}^m\sum_{t=1}^\infty Y_{t,i}\right)Cov\left(\sum_{i=1}^m\sum_{t=1}^\infty Y_{t,i},\left(\sum_{i=1}^m\sum_{t=1}^\infty tY_{t,i}\right)^3\right)}{E\left(\sum_{i=1}^m\sum_{t=1}^\infty tY_{t,i}\right)^5}\\
&= \quad -\frac{2(1-\rho)^8}{\rho^3(\sum_{i=1}^m\lambda_{0,i})^3}\Bigg(\sum_{i=1}^m\lambda_{0,i}^3\Bigg(\sum_{t=1}^\infty 3t^3\rho^{3t}+18\sum_{t=1}^\infty\sum_{\substack{u=1\\u\neq t}}^\infty\sum_{\substack{u=1\\v\neq t,u}}^\infty tuv\rho^{t+u+v}\Bigg)\\
&\quad+\sum_{i=1}^m\lambda_{0,i}^2\Bigg(6\sum_{t=1}^\infty i^3\rho^{2t}+6\sum_{t=1}^\infty\sum_{\substack{u=1\\u\neq t}}^\infty t^2u\rho^{t+u}\Bigg)+\sum_{i=1}^m\lambda_{0,i}\Bigg(\sum_{t=1}^\infty t^3\rho^t\Bigg)\\
&\quad+\sum_{i=1}^m\sum_{\substack{l=1\\l\neq i}}^m\lambda_{0,i}^3\lambda_{0,l}\Bigg(6\sum_{t=1}^\infty\sum_{\substack{u=1\\u\neq t}}^\infty t^2u\rho^{3t+u}\Bigg)+\sum_{i=1}^m\sum_{\substack{l=1\\l\neq i}}^m\lambda_{0,i}^2\lambda_{0,l}\Bigg(15\sum_{t=1}^\infty t^3\rho^{3t}\\
&\quad+36\sum_{t=1}^\infty\sum_{\substack{u=1\\u\neq t}}^\infty t^2u\rho^{2t+u}+18\sum_{t=1}^\infty\sum_{\substack{u=1\\u\neq t}}^\infty\sum_{\substack{v=1\\v\neq t,u}}^\infty tuv\rho^{t+u+v}\Bigg)\\
&\quad+\sum_{i=1}^m\sum_{\substack{l=1\\l\neq i}}^m\lambda_{0,i}\lambda_{0,l}\Bigg(6\sum_{t=1}^\infty t^3\rho^{2t}+3\sum_{t=1}^\infty\sum_{\substack{u=1\\u\neq t}}^\infty t^2u\rho^{t+u}\Bigg)\\
&\quad+\sum_{i=1}^m\sum_{\substack{l=1\\l\neq i}}^m\lambda_{0,i}^2\lambda_{0,l}^2\Bigg(3\sum_{t=1}^\infty t^3\rho^{4t}\Bigg)+\sum_{i=1}^m\sum_{\substack{l=1\\l\neq i}}^m\lambda_{0,l}^2\lambda_{0,i}\Bigg(3\sum_{t=1}^\infty t^3\rho^{3t}\Bigg)\\
&\quad+\sum_{i=1}^m\sum_{\substack{l=1\\l\neq i}}^m\sum_{\substack{r=1\\r\neq i,l}}^m\lambda_{0,i}\lambda_{0,l}\lambda_{0,r}\Bigg(18\sum_{t=1}^\infty t^3\rho^{3t}+18\sum_{t=1}^\infty\sum_{\substack{u=1\\u\neq t}}^\infty t^2u\rho^{2t+u}\\
&\quad+6\sum_{t=1}^\infty\sum_{\substack{u=1\\u\neq t}}^\infty\sum_{\substack{v=1\\v\neq t,u}}^\infty tuv\rho^{t+u+v}\Bigg)\Bigg),
\end{aligned}
\tag{2.E.7}
$$

$$
\begin{aligned}
\frac{\mu_4(\sum_{i=1}^m\sum_{t=1}^\infty tY_{t,i})\mathbb{E}(\sum_{i=1}^m\sum_{t=1}^\infty Y_{t,i})}{2\mathbb{E}(\sum_{i=1}^m\sum_{t=1}^\infty tY_{t,i})^3} &= \frac{1+11\rho+11\rho^2+\rho^3}{\rho(\sum_{i=1}^m\lambda_{0,i})}\\
&\quad+\frac{3\sum_{i=1}^m\lambda_{0,i}^2(1+11\rho^2+11\rho^4+\rho^6)}{(\sum_{i=1}^m\lambda_{0,i})^2}.
\end{aligned}
\tag{2.E.8}
$$

The approximate variance of $\hat{\rho}$ in the multiple series setting is given as the sum of (2.E.1), (2.E.2), (2.E.3), (2.E.4), (2.E.5), (2.E.6), (2.E.7) and (2.E.8).

2.F Mean, Bias and Variance of Numerically Estimated Parameters from the Simulation

Table 2.F.1: Mean of Estimates of λ_0 Using PD (10,000 Simulated Demand Series for each Combination of Parameters)

	$\rho = 0.97$			$\rho = 0.98$			$\rho = 0.99$		
n	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$
12	10.247	50.144	100.152	10.324	50.186	100.049	10.471	50.495	100.407
24	10.036	50.017	100.024	10.034	50.049	100.094	10.110	50.022	99.939
60	10.021	50.013	100.014	10.021	50.007	100.008	10.022	50.017	99.967

Table 2.F.2: Bias of Estimates of λ_0 Using PD (10,000 Simulated Demand Series for each Combination of Parameters)

	$\rho = 0.97$			$\rho = 0.98$			$\rho = 0.99$		
n	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$
12	0.247	0.144	0.152	0.324	0.186	0.049	0.471	0.495	0.407
24	0.036	0.017	0.024	0.034	0.049	0.094	0.110	0.022	-0.061
60	0.021	0.013	0.014	0.021	0.007	0.008	0.022	0.017	-0.033

Table 2.F.3: Variance of Estimates of λ_0 Using PD (10,000 Simulated Demand Series for each Combination of Parameters)

	$\rho = 0.97$			$\rho = 0.98$			$\rho = 0.99$		
n	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$
12	3.765	20.406	42.538	3.244	19.574	39.734	2.733	15.082	33.267
24	2.156	10.661	21.091	1.949	9.946	19.718	1.630	9.189	18.681
60	1.059	5.263	10.321	0.944	4.507	9.182	0.802	3.858	7.852

Table 2.F.4: Mean of Estimates of ρ Using PD (10,000 Simulated Demand Series for each Combination of Parameters)

	$\rho = 0.97$			$\rho = 0.98$			$\rho = 0.99$		
n	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$
12	0.9676	0.9699	0.9700	0.9761	0.9798	0.9801	0.9836	0.9887	0.9895
24	0.9700	0.9700	0.9700	0.9801	0.9800	0.9799	0.9893	0.9900	0.9901
60	0.9700	0.9700	0.9700	0.9800	0.9800	0.9800	0.9900	0.9900	0.9900

Table 2.F.5: Bias of Estimates of ρ Using PD (10,000 Simulated Demand Series for each Combination of Parameters)

	$\rho = 0.97$			$\rho = 0.98$			$\rho = 0.99$		
n	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$
12	-0.0024	-0.0001	0.0000	-0.0039	-0.0002	0.0001	-0.0064	-0.0013	-0.0005
24	0.0000	0.0000	0.0000	0.0001	0.0000	-0.0001	-0.0007	0.0000	0.0001
60	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 2.F.6: Variance of Estimates of ρ Using PD (10,000 Simulated Demand Series for each Combination of Parameters)

	$\rho = 0.97$			$\rho = 0.98$			$\rho = 0.99$		
n	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$	$\lambda_0 = 10$	$\lambda_0 = 50$	$\lambda_0 = 100$
12	0.0006	0.0002	0.0001	0.0005	0.0001	0.0001	0.0004	0.0001	0.0001
24	0.0001	0.0000	0.0000	0.0001	0.0000	0.0000	0.0001	0.0000	0.0000
60	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Chapter 3

Demand Forecasting at Every Phase of the Life Cycle

3.1 Introduction

Demand forecasts are crucial for many strategic and operational decisions concerning resources, logistics, manufacturing, financial planning and inventory throughout a product's life cycle. Demand patterns may change over the course of a product's life cycle, which is important to capture in forecasting. For example, in spare part management, the product life cycle is typically divided into three phases: growth, maturity and end-of-life (Fortuin, 1980; Dekker et al., 2013), where demand increases during growth, stabilizes during maturity and then declines during the end-of-life phase.

One of the main challenges with forecasting demand over a product's life cycle arises when the decision maker can only observe demand at the early phases of the life cycle and does not have any further information about demand later on. In some cases, historical data is restricted to before a peak demand level or a declining pattern in demand has been identified. This is particularly challenging in the case of long-term strategic decision-making, where demand needs to be forecasted for many periods ahead and potentially cover the remainder of the life cycle. For example, an inventory decision common to aftermarket industries is a final order, also known as a 'last time buy'. LTB decisions can be required at any point in the life cycle and must sustain demand over the rest of the service period.

Another challenge may occur if demand is also intermittent, with periods where no demand is recorded. For example, demand for spare parts in the aftermarket is commonly irregular and intermittent (Hasni et al., 2019). Aftermarkets operate across numerous sectors including automotive, consumer electronics and aviation industries (Durugbo, 2020). However, intermittent demand is a more widespread phenomenon. Sillanpää and Liesiö (2018) and Sarlo et al. (2023) highlight the practicality and relevance of intermittent demand forecasting in retail contexts. Intermittent items, although slow-moving, have high prevalence and in some cases can be attributed to 60% of stock investment (Johnston et al., 2003). Much of the current intermittent demand methodology focuses on stationary demand patterns (Croston, 1972; Syntetos and Boylan, 2005). Where trend is considered, existing approaches focus on the final phase of the life cycle where a decline in demand is present (Moore, 1971; Fortuin, 1980).

The research presented in this chapter is motivated by a real-world application from an automotive manufacturer. The automotive manufacturer has to make operational and strategic decisions at different phases of a product's life cycle, for which they require a reliable demand forecast. In Figure 3.1.1, we illustrate examples of monthly demands for different spare parts where the automotive manufacturer was able to collect data for several phases of each product's life cycle over the last nine years. We observe that demand is non-stationary and follows life cycle patterns. There is an initial rise in the demand pattern before the average demand reaches a maximum level, after which we see a sustained decline. We also observe some monthly periods where no demand is recorded.

A common approach to modelling life cycle demand was developed by Bass (1969).

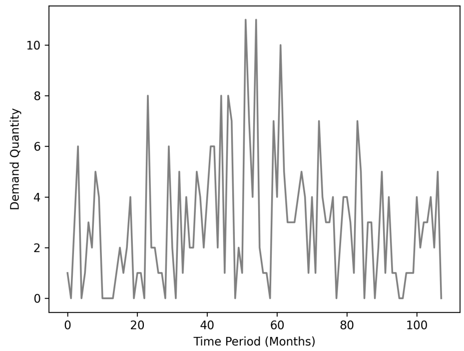

(a) Exhaust Manifold

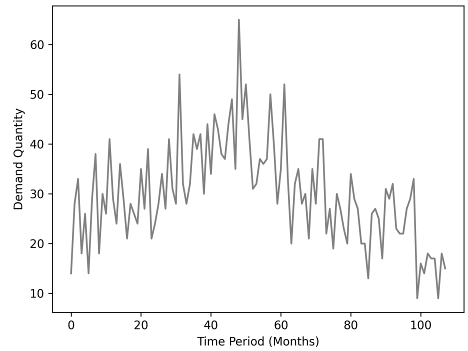

(b) Bracket

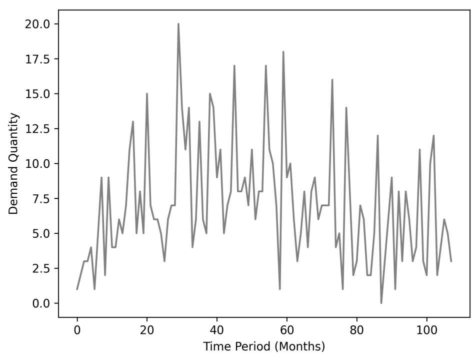

(c) Cover

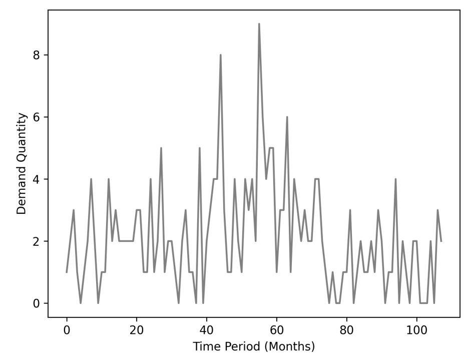

(d) Crossmember

Figure 3.1.1: Demand for SKUs Belonging to an Automotive Manufacturer

Originally designed for demand forecasting of new products, the Bass model is characterised by a bell-shaped curve and consists of a differential equation, explaining the process of customers adopting a new product. Central to the formulation is the assumption that the rate of sales is a linear function of previous sales. The model is favoured for its simplicity whereby solely historical data is used to estimate the model parameters. However, applying the Bass model with limited historical data is known to incur parameter estimation issues (Van den Bulte and Lilien, 1997; Putsis Jr and Srinfvasan, 2000; Meade and Islam, 2006). Implementing the model to forecast demand for products with little or no demand history often relies on utilising sales data of a similar product for which the whole life cycle has been observed.

As the full life cycle for aftermarket SKUs can be several years long, full demand data for similar products is not always available in practice. Manufacturers may not have access to data records far enough into the past to be able to obtain the full demand history of previously sold products. Additionally, technological advancements mean that manufacturers may introduce new products that do not have historical product counterparts. For instance, in the automotive sector, the production of electric vehicles means manufacturers now hold large quantities of stock that do not service standard vehicles.

In this chapter, we contribute methodology to determine long-term forecasts for aftermarket SKUs. We model demand throughout the life cycle as a Poisson process, adopting the Bass curve as a representation of the underlying intensity function. Exploiting the likeness between demand series, we provide a pooling approach to improve forecast accuracy when limited demand history is available. This is done by using the limited demand history for multiple active products to determine the model parameters. Beneficially, we do not rely on complete historical demand series in our implementation. We show that error can be significantly reduced using our multiple series technique for forecasts determined at different phases of the life cycle, including before the height of demand has occurred.

3.2 Literature Review

Life cycle forecasting aims at capturing the pattern of demand throughout all phases of a product's life. The central tenet of life cycle demand forecasting is the Bass model, which has inspired a vast literature (Hauser et al., 2006; Meade and Islam, 2006; Chandrasekaran and Tellis, 2007; Peres et al., 2010; Li and Sui, 2011; Guidolin and Manfredi, 2023). The model is characterised by a simple differential equation which describes product adoption in a market population. By using sufficient sales records of comparable products, Ismail and Abu (2013) apply the Bass model in a real-life example at an automotive company. Sales patterns with related features have also been used when applying the Bass model to determine forecasts for technology products (Ganjeizadeh et al., 2017; Li et al., 2021). Li et al. (2021) forecast demand relying on representative forecasts found using average demand across similar products with life cycle sales records. Representative forecasts are also formed by Hu et al. (2019) by grouping products using time-series-based clustering. The Bass model as well as polynomial, triangle and trapezoid life cycle shapes are used to forecast sales of new computers. Dombi et al. (2018) fit parametric life cycle demand functions to full historical demand series to infer the shape of life cycle curves. Clustering on the parameters of the fitted life cycle demand functions is performed to find typical demand models to forecast the demand for active parts.

Applications of life cycle models have justifiably exploited the likeness between past products to determine forecasts for active products. Life cycle forecasts determined using solely the incomplete demand history of an active part can have low accuracy. For the Bass model, parameter estimation issues when limited historical data is available are well known. Putsis Jr and Srinfvasan (2000) revisit early findings on the estimation of parameters and demonstrate that the peak of demand needs to have been observed to obtain robust and stable estimates (Srinivasan and Mason, 1986). Focusing on the non-linear least squares method, Van den Bulte and Lilien (1997) show that the amount of bias in estimates is in accordance with the degree of truncation of available demand history. Highlighting the issue, Meade and Islam (2006) argue for the relevance of research concerning life cycle modelling for products with limited demand history.

The current methodology to aid long-term decision-making prioritises forecasting demand in the end-of-life phase of the life cycle. Fortuin (1980) design an approach for the LTB problem which considers that demand follows a Gaussian process with an exponentially declining mean. A time-series approach designed by Moore (1971) fits deterioration curves to sales data on a logarithmic scale. In Chapter 2 we introduced a long-term stochastic forecasting model for intermittent SKUs with a declining trend, whereby demand is treated as Poisson. An extension to the model facilitates using demand history from similar products to improve forecasts. These approaches are intended for forecasting during the final phase of the life cycle where a decline in demand is present. As such, demand with a declining pattern needs to be observed to forecast demand over the rest of the life cycle.

Pooling the demand history of multiple products to estimate common model parameters has received increased attention. This approach is based on the assumption that demand for a group of products follows an identical process and is also often referred to as 'cross-learning'. Semenoglou et al. (2021) find that training a single model for multiple SKUs increases forecast accuracy and is much less computationally expensive than individual models for each SKU. In the recent M5 competition, the best-performing approaches utilised joint models using the demand history of multiple products instead of fitting a model for each SKU using solely its own demand information (Makridakis et al., 2022). Cohen et al. (2022) also show that borrowing demand history from other products improves accuracy for newly introduced retail SKUs with insufficient demand history. Lei et al. (2024) propose a pooling framework that integrates both category and product level demand history when determining forecasts for individual SKUs.

The main consideration of the current methodology for intermittent SKUs, such as spare parts, is short-term decision-making. Where long-term circumstances are considered, approaches are concentrated on the end-of-life phase where historical demand data with a declining pattern is used to forecast years ahead. Models which represent multiple phases of the life cycle have yet to focus on intermittent and low-volume items. Additionally, applications of life cycle models, such as the Bass model, to SKUs with little or no demand history rely on similar products for which full life cycle demand is known. We present a forecasting model for spare part SKUs that incorporates the patterns in demand attributed to the product life cycle phases. Additionally, we utilise data pooling techniques and incomplete demand histories of multiple SKUs to forecast cases with insufficient demand history.

3.3 Methodology

In this section, we introduce a long-term stochastic demand forecasting model for aftermarket SKUs that incorporates life cycle patterns. We represent demand in each period throughout the life cycle as a Poisson process with an intensity function that changes with respect to time. We adopt the well-known Bass curve (Bass, 1969) as a representation of the intensity function of the Poisson process. Our initial approach uses the demand history of a single SKU to estimate our model parameters. We develop our methodology to forecast for SKUs with limited demand history by using the combined demand history of multiple products.

3.3.1 The Bass Model

A principle model of life cycle demand forecasting, the Bass model describes how a product is adopted by its market. Two factors are considered with regard to the rate of sales. The first is the probability of a sale at introduction, referred to as the innovation of a product. The second component, the imitation effect, encompasses the impact of previous sales on future sales. Model formulation assumes that the relationship between the rate of sales and the number of previous sales can be described linearly.

Bass (1969) specifies that the cumulative number of sales from time 0 through to time t is

$$S(t) = P\left(\frac{1 - e^{-(p+q)t}}{1 + \frac{q}{p}e^{-(p+q)t}}\right) \tag{3.3.1}$$

where P, $P > 0$, represents the total market size and p, q , $p > 0, q > 0$, are coefficients of innovation and imitation, respectively.

Thus, the number of sales as a function of time is given as

$$s(t) = \frac{dS(t)}{dt} = P\left(\frac{e^{-(p+q)t}p(1+\frac{q}{p})^2}{(1+\frac{q}{p}e^{-(p+q)t})^2}\right). \quad (3.3.2)$$

An acknowledged advantage of the Bass model is its parameter interpretability. Low values of P indicate a small total market size. The innovation coefficient, p, represents the probability of a sale at time $t = 0$. Therefore, high p values correspond to immediate product popularity (Li et al., 2021). When the effect of imitation is larger than the innovation effect ($q > p$) a bell-shaped curve is formed. In the alternate scenario, the peak of sales is observed as soon as a product is put on the market with an onward decline in sales expected. Large values of both p and q are consistent with a product selling very quickly and reaching its maximum market potential. In such cases, sales decay rapidly after observing a peak in sales (Lilien et al., 2017).

3.3.2 Application to Spare Part Demand

The Poisson distribution is a commonly proposed representation of spare part demand (Boylan and Syntetos, 2021). We model demand throughout the life cycle as a Poisson process with an intensity function that changes over time.

For a single unit of a product, we propose that demand incidences follow a Poisson process over the time interval $[0, T]$ with intensity function which follows the Bass curve as introduced in Equation (3.3.2).

Consequently, for all τ such that $0 \leq \delta(\tau-1) < \delta\tau \leq T$, the total demand within time interval $(\delta(\tau-1),\ \delta\tau)$ is a Poisson random variable. Let Y_τ be the number of demand incidences for a single unit of product in the interval $(\delta(\tau-1),\ \delta\tau)$, such that $0 \leq \tau - 1 < \tau \leq T$. We write

$$Y_\tau = Y(\delta(\tau-1), \delta\tau) \sim \text{Poisson}\left(\int_{\delta(\tau-1)}^{\delta\tau} s(t)dt\right).$$

We note that Y_τ is independent of $Y_{\tau+1}$ if $(\delta(\tau-1), \delta\tau) \cap (\delta\tau, \delta(\tau+1)) = \emptyset$. By definition of $s(t)$,

$$\int_{\delta(\tau-1)}^{\delta\tau} s(t)dt = S(\delta\tau) - S(\delta(\tau-1)), \text{ for all } 0 \leq \delta(\tau-1) < \delta\tau \leq T,$$

and

$$\int_0^T s(t)dt = S(T) - S(0) = S(T).$$

For n independent demand observations in consecutive time periods, $y_1, ...y_n$, the resulting likelihood function is given by

$$\begin{aligned} \mathcal{L}(p, q, P; y_1, ..., y_n) &= \prod_{\tau=1}^{n} \left\{ \frac{[S(\tau) - S(\tau-1)]^{y_\tau}}{y_\tau!} e^{-[S(\tau)-S((\tau-1))]} \right\} \\ &\propto \left[\prod_{\tau=1}^{n} [S(\tau) - S(\tau-1)]^{y_\tau} \right] e^{-[S(T)-S(0)]} \\ &= \left[\prod_{\tau=1}^{n} [S(\tau) - S(\tau-1)]^{y_\tau} \right] e^{-S(T)}. \end{aligned}$$

The log-likelihood function is given by

$$\ell(p, q, P; y_1, ..., y_n) = \sum_{\tau=1}^{n} [y_\tau ln(S(\tau) - S(\tau-1))] - S(T). \tag{3.3.3}$$

3.3.3 Extension for Products With Limited Demand History

The issue of poorly estimated parameters when applying the Bass model to products with limited data is well-known (Srinivasan and Mason, 1986; Van den Bulte and Lilien, 1997; Putsis Jr and Srinfvasan, 2000; Meade and Islam, 2006). Additional methodology is needed to determine long-term forecasts for these products. Given the expansiveness of aftermarket inventories, we assume that there exist products with the same life cycle pattern. Therefore, the demand history of SKUs with common life cycle behaviour can be pooled to estimate common model parameters.

Where $Y_{\tau,i} = Y_i(\delta(\tau-1), \delta\tau)$ is the number of demand incidences for a single unit of product i in the interval $(\delta(\tau-1), \delta\tau)$, the demand for product i in period τ is given by

$$Y_{\tau,i} \sim Y_i(\delta(\tau-1), \delta\tau) \sim \text{Poisson}\left(\int_{\delta(\tau-1)}^{\delta\tau} s(t)dt\right)$$

where

$$\int_{\delta(\tau-1)}^{\delta\tau} s(t)dt = S(\delta\tau) - S(\delta(\tau-1)).$$

For n independent demand observations in consecutive time periods belonging to m products, $y_{1,1}, ..., y_{n,1}, ...y_{1,m}..., y_{n,m}$, the resulting likelihood function is given by

$$\begin{aligned}\mathcal{L}(p, q, P; y_{1,1}, ..., y_{n,1}, ...y_{1,m}..., y_{n,m}) &= \prod_{i=1}^{m}\prod_{\tau=1}^{n}\left\{\frac{[S(\tau) - S(\tau-1)]^{y_{\tau,i}}}{y_{\tau,i}!}e^{-[S(\tau)-S(\tau-1)]}\right\} \\ &\propto \left[\prod_{i=1}^{m}\prod_{\tau=1}^{n}[S(\tau) - S(\tau-1)]^{y_{\tau,i}}\right]e^{-[S(T)-S(0)]} \\ &= \left[\prod_{i=1}^{m}\prod_{\tau=1}^{n}[S(\tau) - S(\tau-1)]^{y_{\tau,i}}\right]e^{-S(T)}.\end{aligned}$$

The log-likelihood function is

$$\ell(p, q, P; y_{1,1}, ..., y_{n,1}, ...y_{1,m}..., y_{n,m}) = \sum_{i=1}^{m}\sum_{\tau=1}^{n}[y_{\tau,i} ln(S(\tau) - S(\tau-1))] - S(T). \quad (3.3.4)$$

3.4 Forecast Accuracy

Our forecasting methodology is inspired by the forecasting requirements of long-term operational decisions, such as the LTB decision. When evaluating the performance of a forecasting technique for the LTB, we are concerned with measuring the forecast error between the total true demand and the total forecasted demand over the remaining service period, rather than the forecast error in each period.

We define $f_{t,i}$ as the forecasted demand for series i in period t and F_i as the total forecasted demand across the out-of-sample length, $n+1, ..., N$, where $F_i = \sum_{t=n+1}^{N} f_{t,i}$. Additionally, where $y_{t,i}$ is the demand for series i in period t we specify that the total demand across the out-of-sample period be given by D_i, where $D_i = \sum_{t=n+1}^{N} y_{t,i}$.

The mean percentage error and mean absolute percentage error are scale-independent metrics and thus suitable for measuring performance across a number of demand series that vary in demand size. Over M series, the MPE and MAPE are given as

$$MPE = \frac{100\%}{M} \sum_{i=1}^{M} \left(\frac{D_i - F_i}{D_i} \right)$$

and

$$MAPE = \frac{100\%}{M} \sum_{i=1}^{M} \left| \frac{D_i - F_i}{D_i} \right| .$$

Our methodology can also be used for other long-term managerial decisions for which it may be useful to evaluate forecast performance in every period. To do so, we require a measure that is suitable for non-stationary demand series with periods of zero demand. Accuracy metrics that optimise for the median are not suitable for the series with periods of zero demand as the median is often close to zero. Therefore, these types of measures reward approaches that forecast very low levels of demand (Kolassa, 2016). Measures that involve dividing the forecast error by the true demand in an individual period are also rendered unsuitable due to the periods of zero demand in intermittent series.

The root mean squared scaled error (RMSSE) (Makridakis et al., 2022) scales the mean squared error of the forecast by the mean squared error of a one-step naive forecast on the in-sample demand. The RMSSE for series i is given by

$$RMSSE_i = \sqrt{\frac{\frac{1}{N-(n+1)}\sum_{t=n+1}^{N}(y_{t,i} - f_{t,i})^2}{\frac{1}{n-1}\sum_{t=2}^{n}(y_{t,i} - y_{t-1,i})^2}}.$$

A RMSSE value of less than 1 indicates that the performance of the model is better than the naive forecasting method. Conversely, a RMSSE greater than 1 means that the performance of the model is not better than the naive forecast method. We evaluate performance across M series by calculating the RMSSE for each series and taking the mean across these values. That is,

$$\text{Average RMSSE} = \frac{1}{M}\sum_{i=1}^{M} RMSSE_i.$$

3.5 Simulation

By conducting a simulation, we seek to quantify the accuracy of forecasts determined using our methodology and determine the impact of volume size, in-sample length and the combined innovation and imitation effects on forecast accuracy. We also seek to assess the efficacy of our extension on improving forecasts for products with limited historical data.

3.5.1 Experiment Design

In our experiment, we consider monthly demand over a 10 year period. Three life cycle curves are generated using Equation (3.3.2) with specified parameters, p and q. For given parameters of innovation, p, and imitation, q, the time period of the peak of demand is given by

$$t^* = \frac{ln(q) - ln(p)}{p+q}$$

(Bass, 1969).

We select parameters p and q such that the peak is fixed at 5 years, i.e., $t^* = 60$. Values chosen for p are $0.002, 0.003$ and 0.004. Corresponding values for q are $0.0524, 0.0403$ and 0.0290, respectively. Henceforth, we refer to the curves generated using the parameter pairs as curves A, B and C. Three values for the total market potential, P, are chosen to incorporate the variety of product volume sizes. Table 3.5.1 reports the curve names, corresponding parameters and height of the curves for the chosen values for total market potential. Figure 3.5.1 shows curves A, B and C for monthly periods over a ten-year life cycle for $P = 1000$.

Table 3.5.1: Parameters and Peak Values for Curves A, B and C

			Peak Value		
Curve	p	q	P=500	P=1000	P=5000
A	0.002	0.0524	7.06	14.12	70.60
B	0.003	0.0403	5.82	11.63	58.15
C	0.004	0.0290	4.69	9.39	46.94

For each value of the total market potential, P, and curve A, B and C (parameter pair p, q), we draw 10,000 demand series of length 120 (10 years of monthly periods) from a Poisson distribution which has a mean that changes over time according to the methodology presented in Section 3.3.2. That is, the first demand value is drawn from a Poisson distribution with mean $\int_0^1 s(t) = S(1) - S(0) = S(1)$, the second demand value

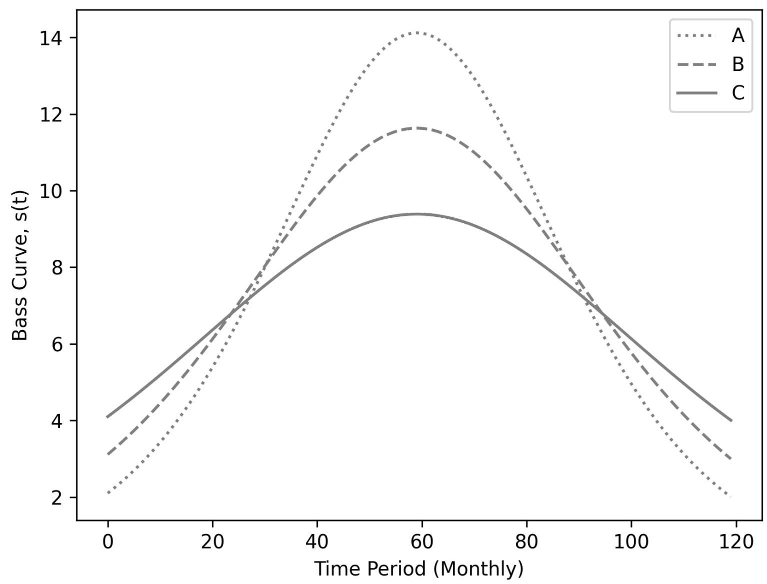

Figure 3.5.1: Life Cycle Curves A, B and C for $P = 1000$

is drawn from a Poisson distribution with mean $\int_1^2 s(t) = S(2) - S(1)$, the third demand value is drawn from a Poisson distribution with mean $\int_2^3 s(t) = S(3) - S(2)$ and so on, for $S(t)$ and $s(t)$ given by Equations (3.3.1) and (3.3.2), respectively, and corresponding parameters p, q, P. Examples of simulated demand series for the different values of P and curves A, B, C are given in Figure 3.5.2. Each subfigure shows a randomly generated demand series for specified P and curve A, B or C.

The in-sample demand for each of the 10,000 demand series is used to estimate the model parameters by numerically maximising the log-likelihood function given by Equation (3.3.3). We select three in-sample lengths in our experiment, $n = 48, 60, 72$. When $p < q$, as is the case for curves A, B and C, the forecasting challenge with our chosen in-sample lengths can be described as follows:

- **Before the peak of demand ($n = 48$)**: The in-sample data ceases one year before the peak of demand occurs. The majority of the life cycle is yet to be observed and forecasts need to capture the incline, peak and decline in demand.

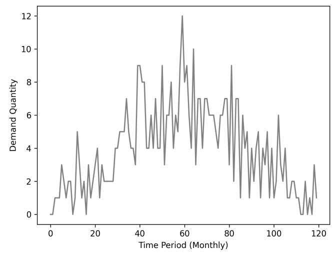

(a) Curve A (P=500)

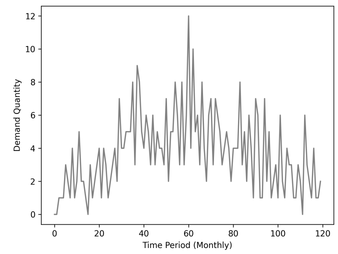

(b) Curve B (P=500)

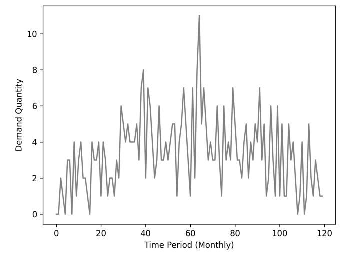

(c) Curve C (P=500)

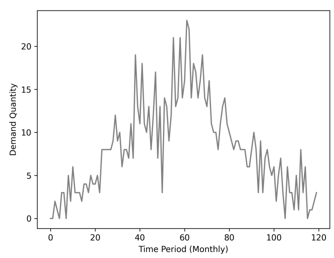

(d) Curve A (P=1000)

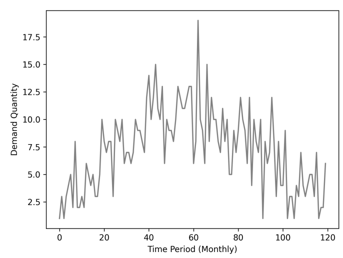

(e) Curve B (P=1000)

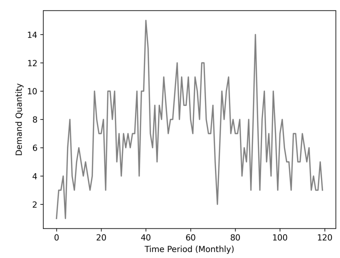

(f) Curve C (P=1000)

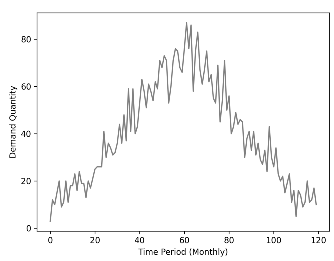

(g) Curve A (P=5000)

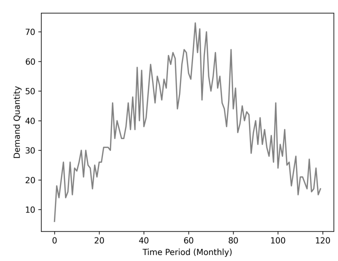

(h) Curve B (P=5000)

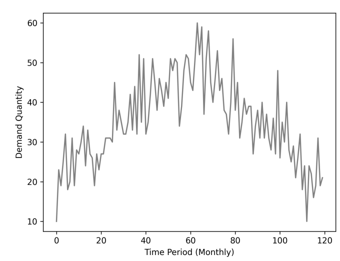

(i) Curve C (P=5000)

Figure 3.5.2: Simulated Demand Series with Intensity Function Given By Curves A, B and C

- **At the peak of demand ($n = 60$)**: Demand will decline over the rest of the planning horizon however the in-sample does not include demand history with a declining pattern.

- **Past the peak of demand ($n = 72$)**: Demand with a declining pattern is limited to 12 periods. Although past the peak of demand, the rest of the life cycle still spans a number of years.

We also test the methodology given in Section 3.3.3. For each of the 10,000 series of random observations of length 120 drawn for each combination of P and curves A, B and C (parameter pair p, q), we draw additional demand series of random observations of length n with the same parameter values. The in-sample demand for each of the 10,000 generated demand series and its additional series are used to estimate the model parameters by numerically maximising the log-likelihood function given by Equation (3.3.4). We consider values of $2, 5, 10$ and 50 for the total number of demand series (1 plus the number of additional series) used in this approach.

3.5.2 Results

We use ‘PBass’ (short for Poisson-Bass) to refer to when only the in-sample from a single series has been used to estimate the model parameters, by numerically maximising Equation (3.3.3). Alternatively, ‘PBassM’ corresponds to our multiple series approach. Parameters are estimated using the in-sample demand of multiple series through the numerical maximisation of Equation (3.3.4). Forecasts are evaluated over the out-of-sample using the metrics outlined in Section 3.4. The mean percentage error and the mean absolute percentage error are presented in Tables 3.5.2 and 3.5.3, respectively. The total number of demand series (1 plus the number of additional series) used in the implementation of PBassM is given in parenthesis. Results pertaining to the root mean squared scaled error, as described in Section 3.4, are presented in Appendix 3.A.

Forecasts determined using the single series approach with data a year past the peak give reasonable accuracy in all cases. In all cases where $n = 72$, we see MAPEs less than 40%. Forecast performance is also good when the series is high volume. When $P = 5000$, the mean absolute percentage error is less than 22% in cases when demand has reached its peak even though no declining demand has yet been observed ($n = 60$).

As expected, accuracy for forecasts generated using data before the peak are impacted by parameter estimation issues. When the demand history is right-truncated before the peak ($n = 48$), accuracy is hardest to achieve for demand series generated

Table 3.5.2: Mean Percentage Error Over 10,000 Generated Demand Series

n	Curve	P	PBass	PBassM (2)	PBassM (5)	PBassM (10)	PBassM (50)
48	A	500	-196.70	-148.01	-81.71	-46.41	-14.74
		1000	-140.21	-89.23	-35.95	-15.91	-5.27
		5000	-33.81	-12.65	-4.34	-2.62	-1.16
	B	500	-91.22	-76.95	-51.59	-33.07	-12.55
		1000	-73.85	-54.90	-29.75	-15.47	-5.37
		5000	-25.22	-12.22	-4.62	-2.56	-0.95
	C	500	-37.07	-35.52	-30.15	-23.35	-11.07
		1000	-34.16	-28.97	-22.26	-14.88	-6.57
		5000	-18.30	-11.65	-5.27	-2.88	-1.04
60	A	500	-50.18	-23.90	-8.62	-5.22	-2.33
		1000	-23.76	-10.36	-4.54	-2.82	-1.84
		5000	-4.17	-2.85	-1.82	-1.56	-1.29
	B	500	-38.98	-25.41	-10.97	-5.98	-2.98
		1000	-24.06	-12.31	-5.06	-3.01	-1.50
		5000	-4.72	-2.46	-1.53	-1.23	-1.01
	C	500	-22.83	-19.64	-12.70	-8.04	-3.58
		1000	-18.79	-13.49	-7.18	-4.02	-1.87
		5000	-5.76	-2.75	-1.39	-1.01	-0.73
72	A	500	-10.96	-6.13	-3.77	-3.07	-2.41
		1000	-5.41	-3.66	-2.55	-2.19	-1.94
		5000	-2.37	-2.11	-1.93	-1.85	-1.74
	B	500	-13.60	-7.63	-3.68	-2.58	-1.99
		1000	-7.32	-4.43	-2.56	-2.06	-1.56
		5000	-2.12	-1.53	-1.33	-1.28	-1.24
	C	500	-13.12	-8.78	-4.80	-3.21	-1.93
		1000	-8.23	-4.69	-2.49	-1.76	-1.22
		5000	-2.20	-1.38	-1.09	-0.98	-0.84

Table 3.5.3: Mean Absolute Percentage Error Over 10,000 Generated Demand Series

n	Curve	P	PBass	PBassM (2)	PBassM (5)	PBassM (10)	PBassM (50)
48	A	500	246.57	188.37	111.55	69.56	27.32
		1000	181.25	121.67	59.15	34.03	14.51
		5000	57.33	30.87	16.46	11.55	5.33
	B	500	141.46	118.17	82.21	57.29	25.20
		1000	114.66	87.89	54.07	34.14	15.01
		5000	49.98	30.76	17.31	11.90	5.29
	C	500	89.38	78.22	62.18	49.08	24.65
		1000	77.69	64.22	47.80	34.72	16.87
		5000	44.46	31.57	18.88	13.10	5.95
60	A	500	81.97	48.10	25.37	17.74	8.98
		1000	47.87	28.71	16.85	11.93	6.36
		5000	16.43	11.40	7.19	5.32	2.90
	B	500	73.19	51.53	29.74	20.09	10.36
		1000	50.64	32.53	19.00	13.16	6.75
		5000	18.37	12.48	7.94	5.72	3.01
	C	500	60.21	48.94	33.18	23.48	11.92
		1000	47.49	35.57	22.43	15.64	7.93
		5000	21.25	14.29	8.97	6.49	3.28
72	A	500	32.02	21.71	14.34	10.99	7.57
		1000	20.82	14.96	10.05	7.82	5.42
		5000	9.06	6.62	4.62	3.69	2.72
	B	500	37.29	25.26	15.92	11.96	7.92
		1000	24.70	17.24	11.19	8.48	5.59
		5000	10.28	7.34	4.89	3.79	2.59
	C	500	39.90	29.20	18.96	14.02	8.60
		1000	28.01	19.82	12.76	9.48	5.95
		5000	11.82	8.41	5.51	4.18	2.65

from Curve A. This curve has a larger imitation effect causing a faster acceleration and decay in demand than the other two curves, making the life cycle pattern in demand much harder to estimate.

Forecasts using the single series approach incur more severe errors when demand volume is low. Lower volume series contain more frequent zero demand periods so it is particularly difficult to effectively estimate parameters. In severe cases, the procedure fails to identify a peak of demand across the total 10 year period, resulting in large negative mean percentage errors. However, in every case, the bias in the forecast reduces as the number of series included in PBassM increases.

Using in-complete demand history from series with identical parameters improves forecast accuracy at all considered phases of the life cycle. The multiple series approach successfully compensates forecast accuracy for low-volume series of varying in-sample length, offering improvements to the MAPE in every case. The more series included in the implementation of PBassM the smaller the reported MAPE. In the case where demand history is only available up to 12 months before the peak of demand is due to be observed ($n = 48$), we see a drastic improvement in forecast accuracy with the inclusion of demand history from additional series. PBassM achieves a mean absolute percentage error of less than 30% when 50 demand series are used in its implementation.

For PBass, PBassM (10) and PBass (50), Figure 3.5.3 shows examples of forecasted curves determined using an in-sample length of $n = 60$ and the true curves used to simulate the demand series for a value of $P = 500$. Similar figures for in-sample lengths $n = 48$ and $n = 72$ are given in Appendix 3.B. The forecasted curves are close to the true curves when a large number of series are used in the implementation of PBassM (i.e., for PBassM (50)). The greatest misalignment between the true curves and the forecasted curves is seen when using PBass and for Curve A in particular. Overall, the alignment of the forecasted curves with the true curves improves as the number of series used in the implementation of PBassM increases. These observations are in accordance

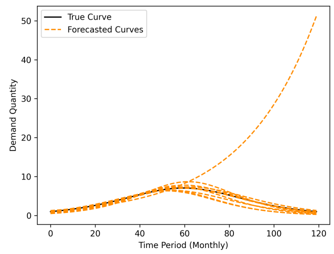

(a) PBass - Curve A

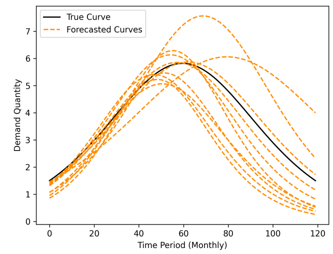

(b) PBass - Curve B

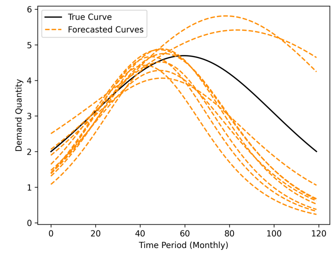

(c) PBass - Curve C

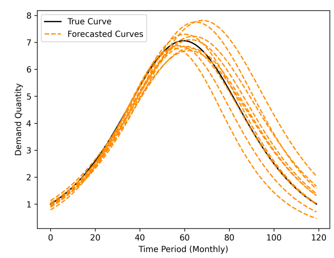

(d) PBassM (10) - Curve A

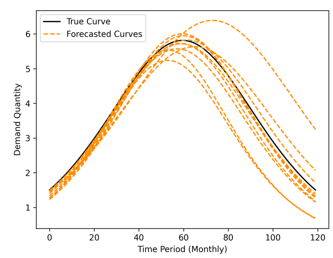

(e) PBassM (10) - Curve B

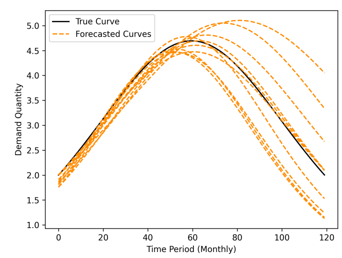

(f) PBassM (10) - Curve C

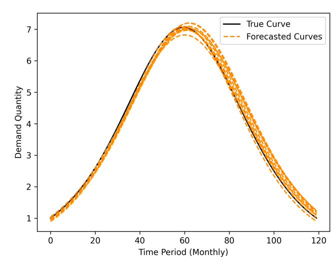

(g) PBassM (50) - Curve A

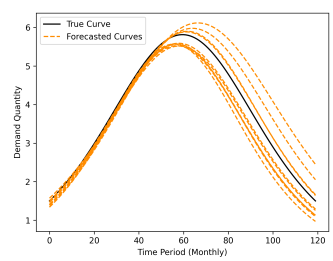

(h) PBassM (50) - Curve B

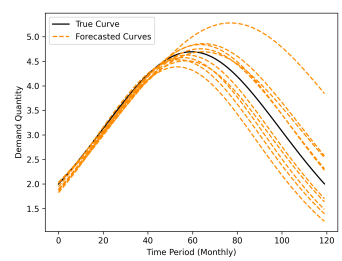

(i) PBassM (50) - Curve C

Figure 3.5.3: Comparison of True Curves and Example Forecasted Curves for $P = 500$ and In-Sample Length of 60 Monthly Time Periods

with the results for the MPE and MAPE given in Tables 3.5.2 and 3.5.3.

The experimentation in this section is done in a controlled setting whereby the demand series are generated based on the assumptions of our methodology. In the following section we test our approach on real data to reflect the challenges faced by practitioners in industry.

3.6 Real Data

We assess performance of our proposed methodology on real data using monthly demand for 175 automotive SKUs obtained from a manufacturer. The demand data covers 9 years of the life cycle, from 2014-2022, consisting of 108 monthly demands per SKU.

In our experiment, we consider in-sample lengths of 48, 60 and 72 monthly periods (4, 5, 6 years). Forecasts are evaluated for the total demand over the remaining time period, corresponding to an out-of-sample of 60, 48 and 36 monthly periods (5, 4 and 3 years), respectively. We estimate the parameters for PBass by numerically maximising Equation (3.3.3).

We apply the multiple series approach, PBassM, using the in-sample demand for all of the SKUs in the dataset by numerically maximising Equation (3.3.4). Before the parameter estimation, the demand for each SKU is standardised by dividing the demand in each period in the in-sample by the total demand across its in-sample length. After common parameters have been found, forecasts are de-standardised by multiplying the forecasted demand for each SKU by the total demand across its in-sample length.

Table 3.6.1: Forecast Performance Across 175 Automotive SKUs

	MPE		MAPE		Average RMSSE	
n	PBass	PBassM	PBass	PBassM	PBass	PBassM
48	-13.71	-28.22	58.28	31.60	1.24	0.94
60	-16.57	-19.01	51.77	25.02	0.99	0.79
72	-3.94	-0.78	38.32	21.55	0.82	0.74

In line with earlier findings, forecast errors for PBass are large when limited historical data is available, especially when the peak of demand has yet to be observed. Table 3.6.1 reports the metrics outlined in Section 3.4 for both approaches calculated across the 175 SKUs. We see an improvement in forecast accuracy as the in-sample length increases. PBass achieves a MAPE less than 40% when $n = 72$ despite a declining pattern of demand having been observed for a short period of time. This showcases the

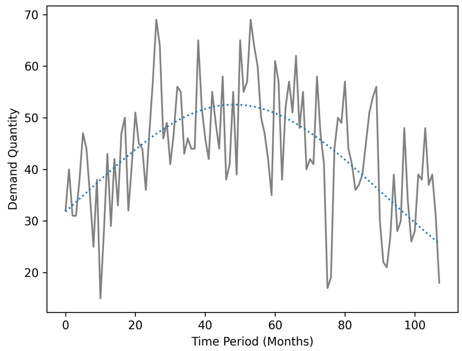

(a) Exhaust Heat Shield

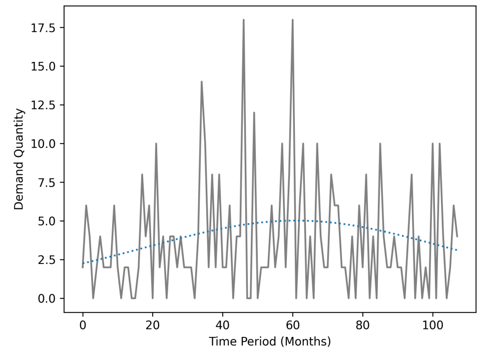

(b) Valve

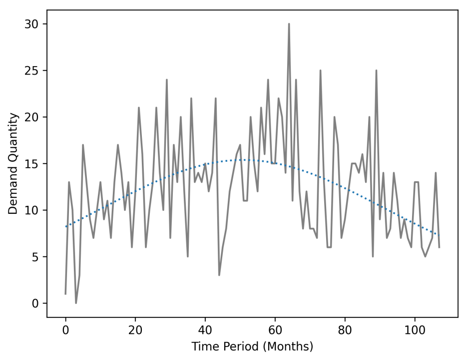

(c) Seal

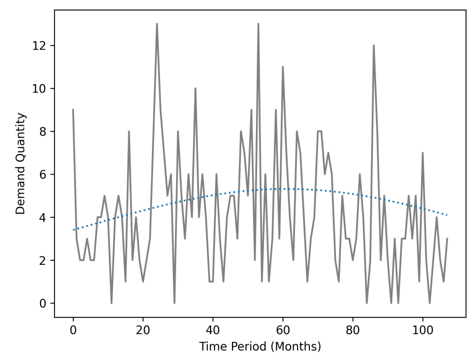

(d) Pulley

Figure 3.6.1: Examples of Forecasts Determined Using PBass with an In-Sample Length of 72 Monthly Time Periods

strength of our single series methodology for long-term decision-making in the end-of-life phase. Example forecasts using PBass with in-sample length $n = 72$ are shown in Figure 3.6.1.

Utilising series with similar life cycle behaviour using our methodology is effective on real data, even when available demand history is limited. PBassM produces more accurate forecasts than PBass and acquires a mean absolute percentage error 47% lower on average. Figure 3.6.2 shows the percentage errors and absolute percentage errors for PBass and PBassM with varying lengths of demand data in the in-sample. The multiple series approach improves upon forecasts determined using the single series, with noticeable improvement in accuracy across all lengths of in-sample considered.

Figure 3.6.3 shows example forecasts determined using PBassM and a short amount of demand history, $n = 48$.

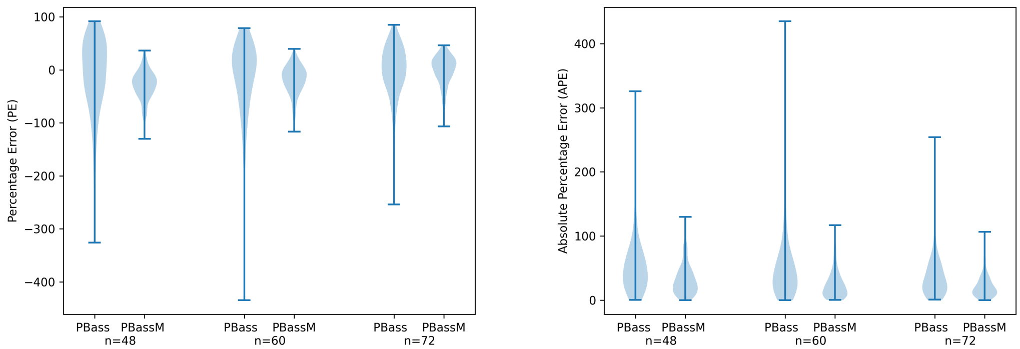

(a) Percentage Error (b) Absolute Percentage Error

Figure 3.6.2: Distribution of Forecast Performance Across the 175 Automotive SKUs

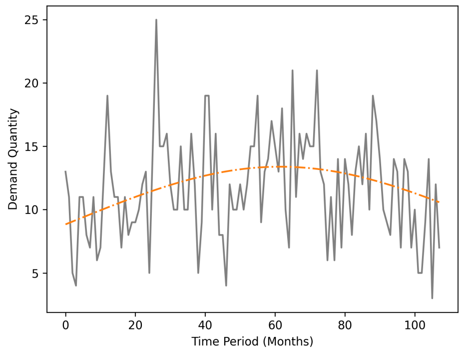

(a) Gasket

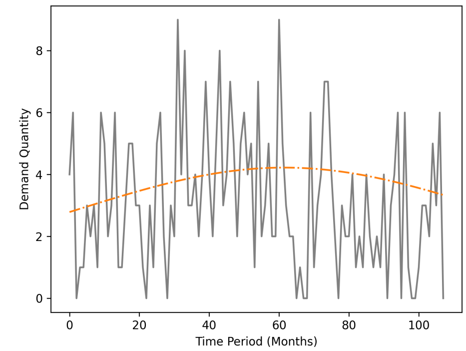

(b) Fuse Box

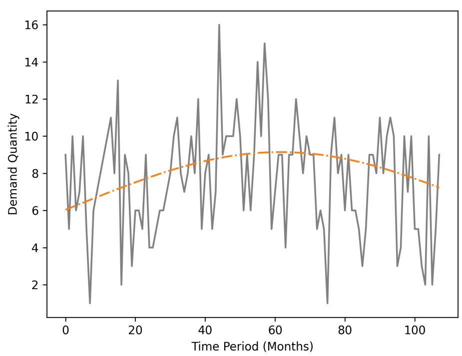

(c) Water Hose

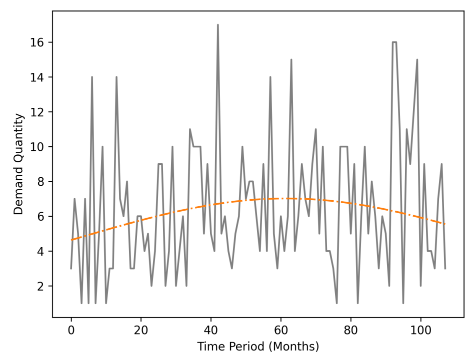

(d) Radiator Hose

Figure 3.6.3: Examples of Forecasts Determined Using PBassM with an In-Sample Length of 48 Monthly Time Periods

3.7 Conclusion

We have introduced a stochastic demand forecasting model for aftermarket SKUs to support long-term decision-making. We model demand over time as a Poisson process with a non-stationary intensity function that follows the well-established Bass curve (Bass, 1969). Acknowledging understood parameter estimation issues with the Bass model when historical demand is right-truncated we further develop our methodology to improve demand forecasts. Our extension pools the incomplete demand history of other products with similar life cycle patterns to estimate model parameters. We obtain the parameters of our model and its extension using maximum likelihood estimation.

In numerical experiments we find that forecast accuracy improves as the amount of historical data and the size of demand increases. Forecasts determined using our model extension have improved accuracy compared to our initial approach and forecast error reduces as the number of series used in the multiple series extension increases.

We applied our methodology to a dataset consisting of demand for 175 aftermarket SKUs obtained from an automotive manufacturer. We find that our initial approach performs well on the real data for larger samples of demand history and when the peak of demand is likely to have passed. This is useful for long-term forecasting for SKUs in the end-of-life phase, particularly in cases where only a small amount of demand with a declining pattern has been observed. Our model extension greatly improves forecast accuracy for products with a limited demand history. As such, our multiple-series approach is especially useful for constructing long-term forecasts throughout the life cycle, particularly when a declining demand pattern or peak in demand value is yet to be observed.

We estimate the parameters of our model extension using the incomplete demand history of all SKUs in the dataset. In the case where inventories have parts with varying life cycle behaviour, SKUs could first be separated into groups with similar life cycle patterns. Product information could be useful for grouping parts with similar demand patterns together. The selection of accurate grouping attributes would require careful consideration and is likely to be specific to the relevant context. Alternatively, time-series-based clustering approaches which group series based on the overall growth pattern of the series and allow for temporal distortion could be employed.

In this work, we focus on representing demand in each period as a Poisson random variable. A potential area for future research is to explore alternative representations of demand, such as different demand distributions that are useful for spare parts (Syntetos et al., 2012, 2013). Additionally, the underlying pattern for demand incidences is assumed to follow the Bass curve. However, other life cycle shapes may also be beneficial for this problem. We highlight these considerations as possible areas for future research.

3.A Simulation Results (RMSSE)

Table 3.A.1: Average Root Mean Squared Scaled Error Over 10,000 Generated Demand Series

n	Curve	P	PBass	PBassM (2)	PBassM (5)	PBassM (10)	PBassM (50)
48	A	500	6.30	4.86	2.96	2.01	1.21
		1000	6.44	4.31	2.24	1.47	1.04
		5000	4.17	2.30	1.44	1.19	0.94
	B	500	3.23	2.76	2.04	1.57	1.06
		1000	3.61	2.82	1.88	1.36	0.98
		5000	3.30	2.12	1.38	1.12	0.89
	C	500	1.90	1.72	1.48	1.29	0.97
		1000	2.22	1.92	1.55	1.28	0.95
		5000	2.71	2.02	1.37	1.11	0.87
60	A	500	1.76	1.20	0.89	0.82	0.76
		1000	1.46	1.06	0.86	0.80	0.75
		5000	1.17	0.98	0.84	0.80	0.75
	B	500	1.54	1.23	0.93	0.83	0.76
		1000	1.49	1.12	0.88	0.81	0.75
		5000	1.23	1.00	0.86	0.80	0.75
	C	500	1.27	1.14	0.95	0.85	0.77
		1000	1.37	1.15	0.93	0.83	0.76
		5000	1.34	1.06	0.88	0.81	0.75
72	A	500	0.78	0.70	0.66	0.64	0.63
		1000	0.76	0.70	0.66	0.64	0.63
		5000	0.74	0.69	0.66	0.64	0.63
	B	500	0.87	0.76	0.70	0.68	0.66
		1000	0.84	0.75	0.70	0.68	0.66
		5000	0.81	0.74	0.69	0.68	0.66
	C	500	0.92	0.82	0.74	0.71	0.69
		1000	0.92	0.81	0.73	0.71	0.69
		5000	0.88	0.79	0.73	0.71	0.69

3.B Comparison of True Curves and Example Forecasted Curves for $P = 500$ and In-Sample Lengths 48 and 72 Monthly Time Periods

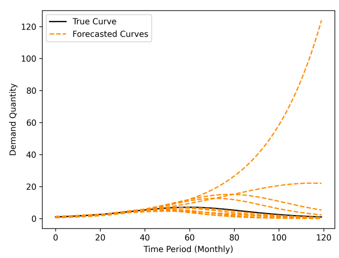

(a) PBass - Curve A

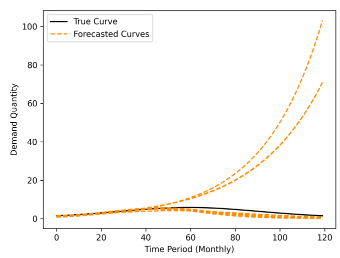

(b) PBass - Curve B

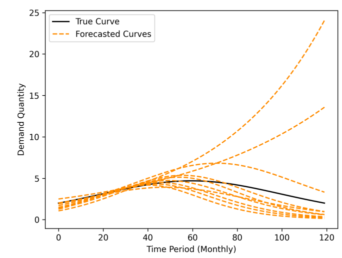

(c) PBass - Curve C

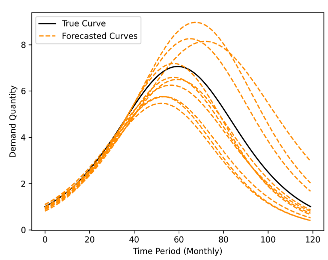

(d) PBassM (10) - Curve A

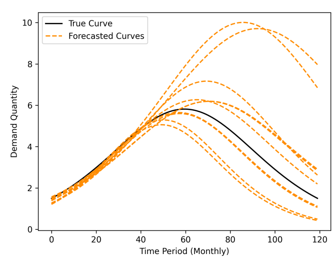

(e) PBassM (10) - Curve B

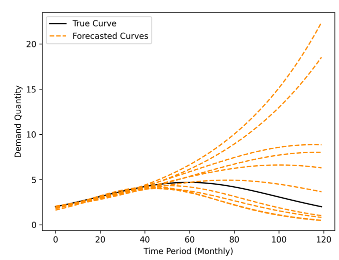

(f) PBassM (10) - Curve C

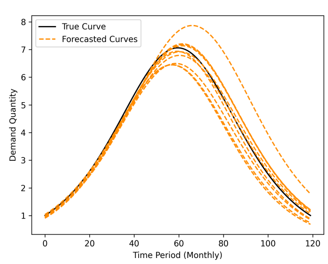

(g) PBassM (50) - Curve A

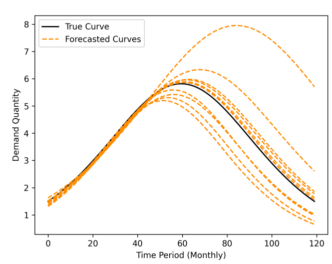

(h) PBassM (50) - Curve B

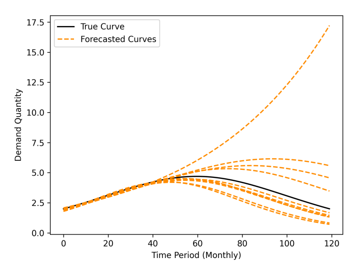

(i) PBassM (50) - Curve C

Figure 3.B.1: Comparison of True Curves and Example Forecasted Curves for $P = 500$ and In-Sample Length of 48 Monthly Time Periods

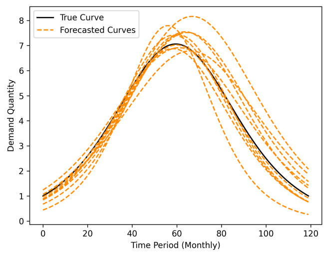

(a) PBass - Curve A

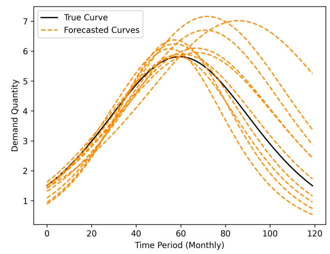

(b) PBass - Curve B

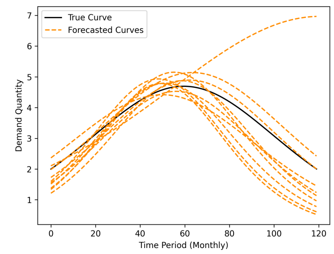

(c) PBass - Curve C

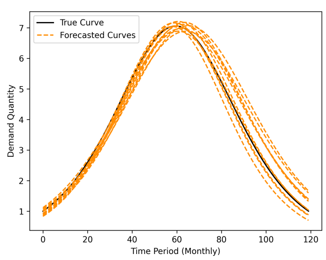

(d) PBassM (10) - Curve A

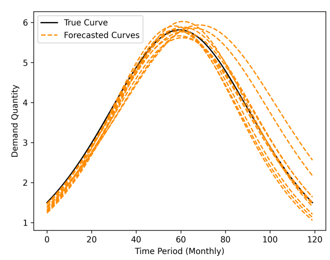

(e) PBassM (10) - Curve B

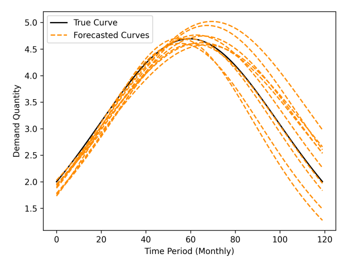

(f) PBassM (10) - Curve C

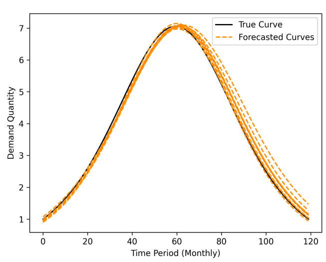

(g) PBassM (50) - Curve A

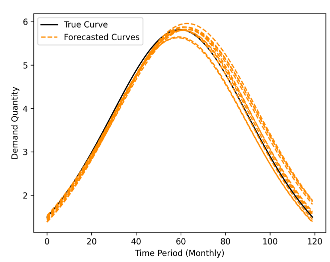

(h) PBassM (50) - Curve B

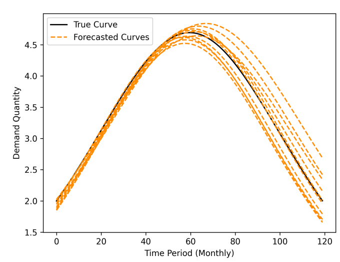

(i) PBassM (50) - Curve C

Figure 3.B.2: Comparison of True Curves and Example Forecasted Curves for $P = 500$ and In-Sample Length of 72 Monthly Time Periods

Chapter 4

Effective Outreach Design for Long-Term Impact

4.1 Introduction

The STEM sectors are paramount to securing future prosperity and tackling social and economic challenges that impact the worldwide population. However, there is a distinct gap between the STEM skills required by industry and the current expertise of workers (Royal Society, 2021; Campaign for Science and Engineering (CaSE), 2023). In England, almost 50% of the working-age population have the level of numeracy expected of a primary school child (Department for Business, Innovation & Skills, 2012). The problem is made worse by the critical shortage of mathematics teachers, where currently at least 12% of maths lessons are taught by non-specialists (Weale, 2023). At UK universities, many mathematics departments are facing a crisis of low uptake, leading to the discontinuation of undergraduate courses (Saunders, 2024). Despite their importance, the STEM sectors face many difficulties in filling necessary roles in industry, charity and government.

The overall shortage of professionals is not the only problem. The STEM community suffers from severe diversity issues (Science and Technology Committee, 2023). The STEM workforce in the UK is predominately comprised of white men, who occupy 65% of roles (British Science Association, 2020). Disabled people are underrepresented within STEM sectors with a larger gap in representation for disabled women than disabled men (British Science Association, 2020). People who belong to the LGBTQ+ community are also underrepresented, with 30% of LGBTQ+ young people specifying concerns over discrimination as the reason they would not choose to pursue a career in STEM (STEM Women, 2021). Students from lower socio-economic backgrounds are also critically underrepresented within STEM sectors (Weale, 2024).

Public dissemination and outreach initiatives help to combat issues within STEM sectors including shortages in the workforce, the underrepresentation of particular demographics and low levels of public numeracy. Outreach activities include talks and workshops which aim to nurture interest in STEM disciplines and encourage future generations to pursue careers in STEM fields. Communicators also work to engage the wider public in STEM-related topics, attracting attention to mathematical ideas and concepts which have tangible relationships to people's everyday lives. Initiatives of this kind work to bridge the gap between the scientific community and the rest of society. In this chapter, the term 'outreach' is used to describe activity and content of any form that promotes interest in the mathematical sciences.

This chapter presents a portfolio of outreach work conducted by the author during their PhD. The work is designed based on core principles and focuses on disseminating techniques within statistics and OR related to forecasting and statistical modelling. A multitude of dissemination types are considered including published articles targeted to audiences with differing levels of scientific knowledge. A talk and an interactive workshop delivered to school students are also presented. A reflection on the impact of this work is given using audience feedback and the author's own experiences of delivery. Motivated by the challenge of overcoming gender disparity within STEM fields, the work presented incorporates attempts to address barriers to mathematical identity. Further, this work puts forward the argument that designing outreach to target obstacles to belonging in mathematics can maximise the long-term impact of initiatives, particularly towards the aim of improving issues of gender equity in the mathematical sciences.

4.2 Background

4.2.1 Public Dissemination and Outreach

Outreach activity has many forms. Talks and workshops tailored to school-age students take place nationwide and are arranged by organisations, such as charities (Mathematics Education Innovation (MEI), 2018; Royal Instituition, 2024) and universities. These activities generate awareness of mathematical topics and career paths. Emphasis is placed on non-traditional content and delivery where interaction is often encouraged. Outreach professionals also seek to reach public audiences. Scientific writing is published in mainstream news outlets (Pagel, 2023; Yates, 2023; Steckles, 2025) and as popular science books (Bellos, 2011; Fry, 2018; Rutherford, 2018; Yates, 2019; Imafidon, 2022). Examples of science television programmes include the BBC series *The Secret Genius of Modern Life*, renewed for a second season in 2024 (Bruce, 2022).

Outreach activity has been found to have a great influence on its audiences. A reflection on experiences of school students, Vennix et al. (2018) found that learning via outreach was perceived positively. A survey produced by Appel et al. (2021) indicated that a single event or individual had positively influenced the choice to pursue a career in STEM for 97% of participants. More specifically, topics within OR have been shown to improve interest and skills related to STEM in programmes designed for school students (Raffaele and Gobbi, 2021; Colajanni et al., 2023). An investigation conducted by Steckles et al. (2020) found that many undergraduates undertaking degrees in the mathematical sciences had engaged in informal mathematical activity before studying at university.

Although highly influential, the effective design of effective outreach is insufficiently studied. This applies especially to outreach initiatives involving topics from the mathematical sciences (Steckles et al., 2020). By reviewing work by STEM organisations, Tillinghast et al. (2020) conclude that outreach activities should 'support and increase the understanding, awareness and interest in STEM disciplines'. Others argue that the aim of fostering enthusiasm runs in tandem with addressing matters of workforce shortages. Johnson and Mulligan (2016) reason that the focal point of outreach design should be the connection of content to careers so as to encourage young people to consider mathematical career pathways. An outcome of outreach activities, as argued by Crawford et al. (2021), is to improve the readiness of participants to fill the gap in STEM professionals and to increase total innovation by especially promoting to underrepresented groups.

Some outreach professionals have designed initiatives solely towards the aim of improving diversity within the mathematical community. The talk *The Black Heroes of Mathematics* aims to shed light on overlooked contributions to the mathematical sciences made by black mathematicians. Delivered at multiple events across the UK and available online, the talk frankly discusses obstacles facing ethnic minorities in STEM subjects (Chamberlain, 2020). Organisation *Lathisms* highlights the contributions of Latinx and Hispanic mathematicians by recognising the achievements of a mathematician for each day of Hispanic Heritage Month. The digital platform *Womanthology* promotes female role models in science through a multitude of dissemination outlets including podcasts and articles (Womanthology, 2024). The creators of the video *Faces of Women in Mathematics* aim to celebrate the work and diversity of female mathematicians (Linke and Hunsicker, 2018). Their work has collated video clips featuring almost 250 women from over 30 countries who are active in mathematics.

4.2.2 Gender Disparity in the Mathematical Sciences

There is a long-standing issue concerning gender representation within STEM sectors. Women are underrepresented in all career levels within the STEM workforce (British Science Association, 2020) and hold just 24% of core positions (WISE, 2019). There is no contention that the gender imbalance extends to academia. The progressive disparity between the number of men and women in STEM-related academic fields is colloquially referred to as the 'leaky pipeline'. In mathematics, national benchmarking finds that women are underrepresented at every level from A-level through to Professorship (London Mathematics Society, 2023). At the highest level of academic seniority, the disparity is stark. In the UK, just 14% of professors in the mathematical sciences are female (London Mathematics Society, 2023). Of the top one hundred most-cited academics in OR identified on *Google Scholar*, only one woman is listed (Carroll and Esposito Amideo, 2024).

The gender disparity within STEM in large part arises from gender bias. Gender stereotypes concerning STEM subjects are influenced by societal views and perpetuated often unconsciously (Reinking and Martin, 2018). The notion that boys' interests are better aligned with computer science and engineering disciplines, is found to be upheld by children as young as six years old (Master et al., 2021). Unconscious gender bias can have significant consequences, although judgements occur implicitly. Research into the impact of gender on the experiences of women and gender minorities pursuing PhDs in mathematics found imposter syndrome, a lack of female mentors, differential supervisory treatment and experiences of sexism to be key themes (Reis and Wadkin, 2023). Explicit gender bias results in deliberate acts of prejudice and, in the worst case, severe forms of discrimination and violence. Such forms include sexual harassment, experienced by approximately 10% of women postgraduate researchers in higher education (Bull, 2024). The widespread endorsement of gender stereotypes diminishes girls' aspirations within STEM fields and contributes to discriminatory environments, ultimately leading to poor female representation.

A predominant theme in myths related to gender and STEM subjects relates to competence. The sexist view that women are less capable of scientific pursuits than their male counterparts is far-reaching. Good et al. (2012) find that messages that women are not as talented as men at mathematics and the belief that an individual has a fixed level of mathematical intelligence weakened feelings of belonging for girls in STEM communities. The perception that intelligence is predetermined has also been linked to poor retention of doctoral students across STEM disciplines (Clark et al., 2021). In the same vein, an atypical standard of intelligence is also related to stereotypes associated with mathematical professionals. Shin et al. (2016) and González-Pérez et al. (2020) acknowledge the stereotype that successful scientists are unnaturally gifted in STEM disciplines. This is echoed by Cheryan et al. (2015) who recognise that the inaccurate association of professionals with genius-like qualities curbed girl's interests in STEM subjects. Starr (2018) argue that genius stereotyping of people within the mathematical community erodes the motivation of women to pursue career pathways within the discipline. Bradshaw and Mann (2021) also highlight the negative impact of inaccurate perceptions of competence on underrepresented individuals. They suggest that a lack of transparency from professionals about the inevitably challenging aspect of learning mathematics contributes to this stereotype.

Obstacles to the mathematical community also emerge from stereotypes about the attributes and characteristics of mathematical professionals. Cheryan et al. (2015) argue that beliefs about the culture within computer science and engineering communities contribute to the representation gap between men and women. They found assumptions of masculine interests act as a barrier to belonging for girls in these subjects. Perceptions of unsupportive workspaces and limited collaboration in STEM roles were also found to have a negative impact. Starr (2018) suggests that stereotypes linked to physical appearance and social skills are widely endorsed and lead to non-diverse recruitment.

Stereotypical messaging that perpetuates exclusivity of certain social groups is also discussed by Shin et al. (2016). The affiliation of STEM careers with white males is linked to an adverse sense of belonging for minority individuals in STEM communities.

Solutions to the issue of gender inequality in mathematics are of great importance. González-Pérez et al. (2020) analyse the effect that female role models have on school-aged students. They find that intervention has a positive impact on many factors including girls' enjoyment of mathematics, self-efficacy and wider aspirations in STEM. The authors relate this success to the personal and authentic nature of role model intervention, noting that speakers rely on their own career experiences and journeys in their dissemination. Outreach initiatives led by female role models are also discussed as a solution to tackling stereotypes of identity and competence by Shin et al. (2016). On review of various programmes developed to increase young women's interest in STEM subjects, Sáinz et al. (2022) find that most outreach initiatives were successful in improving students' self-perception about STEM subjects.

4.2.3 Intersection between Gender Disparity and Outreach

Outreach initiatives have been proven to have a positive impact. Definitions of mathematical outreach commonly focus on fostering interest and facilitating the development of skills. This neglects the wider long-term goals to improve the representation of marginalised groups within scientific communities. Initiatives designed to promote diversity within the mathematical community largely operate in isolation, solely aiming towards the goal of improved representation (Lathisms, 2016; Linke and Hunsicker, 2018; Chamberlain, 2020; Womanthology, 2024). Some professionals advocate for outreach that promotes diversity alongside traditional aims. Nevertheless, examples of outreach initiatives that serve multiple objectives are scarce. In discussions about gender imbalances within mathematics, some studies highlight that intervention can help improve perceptions that fuel gender disparity (Shin et al., 2016; González-Pérez et al., 2020; Sáinz et al., 2022). The particular influence of outreach initiatives on female students has also been emphasised (Cronin et al., 2017; Appel et al., 2021). However, there is a shortage of examples of mathematical outreach designed to mitigate factors that negatively influence gender equity.

4.3 Content Design

This section introduces outreach content designed by the author. The content presented focuses on topics in OR, specifically forecasting and modelling. Section 4.3.1 discusses the core principles of dissemination that informed the design of the author's content. Four examples of outreach content are then given. The content takes three formats. Articles are presented in Sections 4.3.2 and 4.3.3. A talk delivered at several outreach events is described in Section 4.3.4. An interactive workshop designed for GCSE-level students is described in Section 4.3.5.

4.3.1 Design Principles

The activity developed by the author was designed based on principles for effective dissemination. To be impactful, outreach activity must be engaging and appealing. Concepts must also be communicated clearly and in appropriate language. Moreover, the most influential outreach activity is relatable to its audience and provides a connection between individuals and the wider mathematical community. These principles in relation to the work by the author are explored below.

Engagement

To address its aims, outreach activity should be engaging. OR methodology taught within a context related to the interests of its audience can improve understanding (Cochran, 2004). Games and puzzles contribute examples that are known to be exciting to students learning OR topics (DePuy and Taylor, 2007; Hartmann, 2018; Arenas-Vasco et al., 2024). More widely, applications to pop culture, including the use of examples from film, television and music can enhance learning experiences (Greenfield, 2007; Chlond, 2016). Applications to sports provide another fruitful example (Cochran, 2004; Goossens and Beliën, 2023; Chan et al., 2024). Playing and watching sports are popular hobbies for many people and data on sporting events is publicly available. This can allow students to apply concepts to non-fictional and familiar scenarios.

Informed by the positive learning outcomes of connecting content to the interests of students, the author's content incorporated applications to popular culture, specifically literature, film, music and television. Applications include Korean pop music (K-Pop) and television dramas (K-Dramas), *William Shakespeare*'s *Hamlet* and *HBO* television hit *Game of Thrones*. The applications reflect mainstream hobbies that are relevant to a wide cross-generational audience. The novelty of these topics was also appealing, given that broader applications are rarely explored within the standard teaching of mathematical subjects and these interests do not conform to the stereotypical hobbies associated with the scientific community. Furthermore, applications of this kind highlight how mathematics can be used to solve problems which have a large economic and cultural impact. The author's choice of subject matter enhanced motivation and engagement, improving the efficacy of content to promote interest in relevant disciplines.

An additional reason for choosing these topics was that much of the information about films, television and music is in the public domain. Utilising the link between OR and authentic applications to inspire learners can be especially effective (Kaiser et al., 2013; Budd, 2015; Spooner et al., 2024; Taranto et al., 2024). The author included genuine data in activity, permitting the grounding of the methodology in a real-world context. At the beginning of in-person activities, the author discussed their PhD research and referred to the real-world application of the models presented in Chapters 2 and 3. For interactive workshops, the author designed content around box office forecasting, motivated by the need for cinemas to recoup losses incurred from the closures during the COVID-19 pandemic. This problem was selected because it was grounded in recent real-life events and had a tangible impact on people's everyday lives. The reliance on authentic problems secures audience interest and provides a solid basis for the dissemination of mathematical ideas.

Communication

Clarity is crucial in the communication of STEM-related topics to non-scientific audiences. The use of technical language can hinder the efficacy of science writing for broader audiences (Rakedzon, 2019). In the design of in-person talks, the author omitted mathematical formulation, focusing instead on using illustrations to communicate core ideas. In an article targeted at a public audience, statistical techniques were explained without mathematical notation. The author used graphs, illustrations and comprehensive descriptions to communicate ideas. For example, when explaining the movement between compartments of the epidemiological SIR (Susceptible, Infectious, Recovered) model (Keeling and Rohani, 2008), a description in lay terms is used in place of the functional forms of the transitions. As relating concepts to tangible examples has been found necessary for effective communication to large audiences (Sullivan Jr et al., 2021), the compartments are renamed to fit within the context of the application. The implementation of the concepts explored in this article were also grounded in application through the use of real data and the incorporation of well-known references, for example, Oscar-winning film *Parasite* and popular television Show *Squid Game*, which enhanced the generality of the writing.

Although the overuse of mathematical formulae is discouraged by experienced communicators, an introduction to terminology can inspire audiences and build STEM skills. When incorporating technical language, speakers can prepare audiences by discussing the meaning of the mathematical formulae before it is presented (Budd, 2015).

In a workshop setting for GCSE students, the author's activity included mathematical notation likely to be unfamiliar to students e.g., the Greek letter sigma (Σ) to denote the sum of terms. Before providing an explanation and an example, the author prepared the audience by acknowledging that the terminology was not expected to be prior knowledge. As peer learning is highly effective in helping students to solidify their understanding (David, 2014), the students were asked to complete tasks in small groups to support the familiarisation of new technical material. These steps ensured the successful incorporation of mathematical notation and formulae.

Consideration of an audience's prior level of mathematical understanding is essential for worthwhile dissemination (Mercer-Mapstone and Kuchel, 2017). The author developed an article on forecasting for mathematics magazine *Mathematics Today*. The audience of the magazine consists of individuals who understand mathematics to the standard expected of graduates in the mathematical sciences. As those reading the magazine have an understanding of mathematical notation and terminology, the descriptions of concepts were accompanied by mathematical formulations. However, given that the target audience encompasses individuals from various backgrounds across industry, academia and teaching, the author was mindful not to presume that the audience would be familiar with forecasting. All methods were explained assuming no prior knowledge of the methodology. The incorporation of mathematical terminology was determined to the level appropriate to the article's audience. The author's understanding of the target audience informed the design of relevant content.

Relatability

Preconceptions surrounding mathematical identity have an adverse effect on the gender gap within the mathematical sciences. Assumptions around mathematical identity include the association of STEM professionals with a particular demographic (e.g. white and male), solely scientific and/or masculine interests and isolated working (Cheryan et al., 2015; Starr, 2018; Shin et al., 2016; Master et al., 2021). Outreach content must be designed with relatability in mind in order to dismantle stereotypes. Multiple studies have found that the visibility of women within STEM careers helps to improve the recruitment of girls in relevant disciplines (Shin et al., 2016; Bertrand and Duflo, 2017; Cassidy et al., 2018; Van Camp et al., 2019). The author initiated verbal activities by introducing themselves to the audience as a mathematician and shared photos of themselves with the other students in their PhD and undergraduate cohorts. This presentation gave a diverse view of individuals pursuing career pathways in OR in terms of gender and race. In doing so the author improved the audience's perceived compatibility with mathematical science professionals.

In improving the relationship between mathematics and the broader public, one technique is to humanise mathematicians by connecting the subject matter to the emotions of the people who study it (Budd, 2015). In a similar vein, role-model intervention can have the most impact when speakers share their unique career journeys and personal experiences (González-Pérez et al., 2020). The in-person activities were designed to incorporate the author's journey in mathematics education. The author discussed setbacks, experiencing low confidence and finding the learning process challenging. In keeping with the attempt to address the inaccurate perception of uncooperative workspaces with STEM fields (Cheryan et al., 2015), the author highlighted a number of individuals who had a positive influence on their career path. This included discussing the collaborative nature of their PhD work and relationships maintained with supervisors, peers and industrial collaborators.

Related to the inaccurate perceptions attributed to the identity of mathematical scientists (Cheryan et al., 2015; Shin et al., 2016; Starr, 2018; Master et al., 2021), the author focused on examples from television, film and music industries that oppose the convention of masculine interests. Whilst exploring the mathematical analysis of the cultural phenomenon, known as the Korean Wave, described as the rapid rise of popularity of Korean cultural exports, the author focused on television and music. The global success of these industries has been driven by their female fan bases (Kim, 2022a; Venema et al., 2023). These examples challenged the conventions of masculine interests of mathematical scientists, removing barriers to mathematical identity. The lack of inclusivity in examples in teaching is also identified as a barrier to mathematical belonging by (Bradshaw and Mann, 2021). To make content more relatable, examples within in-person activity were included to reflect inclusivity. A strategic interaction between two female characters from the television show *Game of Thrones* was given as an example in teaching Game Theory. The inclusion of two female characters was selected to improve the gender diversity within examples, thus enhancing the relatability of resources.

4.3.2 Article: Surfing the Korean Wave

This article uses modelling to analyse the cultural phenomenon known as the Korean Wave. The Korean Wave is described as the rapid rise of popularity of Korean cultural industries including film, television, music, food, beauty and skincare. The phenomenon has attracted global attention and interest, making worldwide headlines. Given that the Korean Wave refers to cultural industries it has relevance to and impacts ordinary people. Subsequently, the topic is pertinent to the general public. The article employs the use of a compartmental model commonly used with epidemiology, the Susceptible, Infectious and Recovered (SIR) model (Keeling and Rohani, 2008), to gain insights into the behaviour of the Korean wave and forecast the height of the phenomenon. A version of the article was published in *Significance Magazine* in December 2023.

Surfing the Korean Wave

In 1994, Kim Young Sam faced a simple statistic. To generate the same global revenue as the American blockbuster *Jurassic Park*, Korea would need to sell more than 1.5 million *Hyundai* cars to foreign customers. For the newly elected President, it was a striking comparison. A product of culture had overshadowed the 'pride of Korea' (Shim, 2006) and what the country was offering to the international market. With the Asian financial crisis looming, Kim Young Sam needed to react. The government established the Cultural Industry Bureau, conglomerates were encouraged to expand business into the media sector and the idea of industrialising Korean culture sent ripples across the country for the very first time (Shim, 2006).

Less than 30 years later, the growing international popularity of Korean popular music (K-pop), television (K-dramas), cosmetics, films, food and other cultural exports is impossible to ignore. Korean boy band *BTS* have ascended worldwide charts, achieving more number one hits on the *Billboard Hot 100* than any other artist this decade (Chin, 2022). The smash hit *Squid Game* became Netflix's most-watched show (Spangler, 2021) in the same year that *Bong Joon-Ho*'s *Parasite* made history as the first non-English language film to win the Oscar for Best Picture. In 2023, Korean girl group *BLACKPINK* headlined one of the world's largest music festivals, *Coachella*, just months after they sold out the UK's *O2* Arena. The meteoric rise of Korea's cultural industries has been dubbed 'Hallyu' and is known in English as 'The Korean Wave'. A special exhibition was recently launched in London's *Victoria & Albert Museum* dedicated to exploring the origins and impact of the Hallyu phenomenon (Kim, 2022b). In the UK and beyond, the Korean Wave certainly has made a splash.

The Hallyu fandom has immense influence. Cultural exports have boosted the Korean economy by more than 12 billion US dollars (Roll, 2021). Though, K-content followers pride themselves on being so much more than just an economic power. At the height of the Black Lives Matter campaign, discriminatory hashtags were overloaded with fan videos of K-pop artists, making racist tweets nearly impossible to find (Kim, 2022a). A million-dollar donation to the movement by *BTS* was matched by the group's fandom, who are known as the *BTS Army* (BBC News, 2020). Later, with a rising number of hate crimes towards the Asian community during the pandemic, fans united in support of the #StopAsianHate campaign, successfully retweeting the message over 1 million times. The original tweet became the most shared of 2021 (Kim, 2022a). Together, the Hallyu fandom makes up an unstoppable force.

A love for K-pop, K-dramas and K-everything-in-between has swept across the globe, inspiring a dedicated group of followers. But just how strong is the Korean Wave? Will passion for K-content keep growing, or will we see a turn of the tide?

Modelling a Kontagion

From the ear-worms characterising K-pop smash hits and the gripping K-drama cliffhangers you *have* to tweet about to the Korean skincare going viral on *TikTok*, Korean cultural exports are infectious. In the age of social media, the sharing of K-content is easier than ever. Algorithms ensure swoonworthy drama scenes shared by one K-drama fan reach the feed of another likely admirer and K-pop enthusiasts tag their friends in uploaded attempts at the latest dance challenges.

SIR models, standing for Susceptible, Infectious and Recovered, are statistical models used to represent infectious processes. They could help us model the Korean Wave, assuming adoration for K-content is contagious and transmission is driven by word-of-mouth and online hype. Of course, there is much more sophisticated statistical analysis that could be done beyond this first approach - when it comes to modelling complex phenomena like Hallyu, where there's a will, there's a wave!

Imagine that we split everyone into three groups defined as:

Fan potentials (S)

This group is made up of potential K-wave fans. If **Fan potentials** come into contact with a supporter of K-content they can become interested in trying it out for themselves.

Kontagious (I)

The **Kontagious** are captivated by K-content. Perhaps these people have recently tried Korean ramyeon or started learning Hangul, the Korean alphabet. Enthralled by their new fascination, this group is engaging with related content online and is capable of sharing their excitement with friends and family.

K-community (R)

K-community members are fully fledged Hallyu fans and members of a fan community. Those in this group could have joined a University K-pop Society, a Korean film club or a Korean skincare group on *Facebook*. Because they have found like-minded friends, they are no longer sharing with others outside of the community. People in this group could have also decided that K-pop, K-dramas or just K-anything really, was not for them. Sad but possible. The key thing is that this group is made up of those who are no longer able to influence **Fan potentials**. They are either in a K-crew or out of one for good.

Let's say that, on average, the number of people one member of the **Kontagious** group interacts with during one time period is the Greek letter, β. Multiplying this by the proportion of the population in the **Fan potentials** group gives us the expected number of people just one **Kontagious** person could reach at this point in time. Multiplying this value by the number of people in the **Kontagious** group tells us the total number of **Fan potentials** who will experience some kind of Korean culture. The number of the **Kontagious** increases with respect to this value at each time step, whereas the size of the **Fan potentials** group reduces by this amount as some become intrigued by K-content.

We can assume that the probability of finding a fan community (or leaving K-content behind) and moving from the **Kontagious** group to the **K-community** is the Greek letter, γ. In a particular time period, the number of **Kontagious** people multiplied by γ tells us how many will become new members of the **K-community**. This value also tells us by how much the **Kontagious** group will shrink.

We now know how the status of the three groups changes over time. An illustration is shown in Figure 4.3.1. By writing this in equational form, we can solve to find a function for the number of individuals in each group with respect to time.

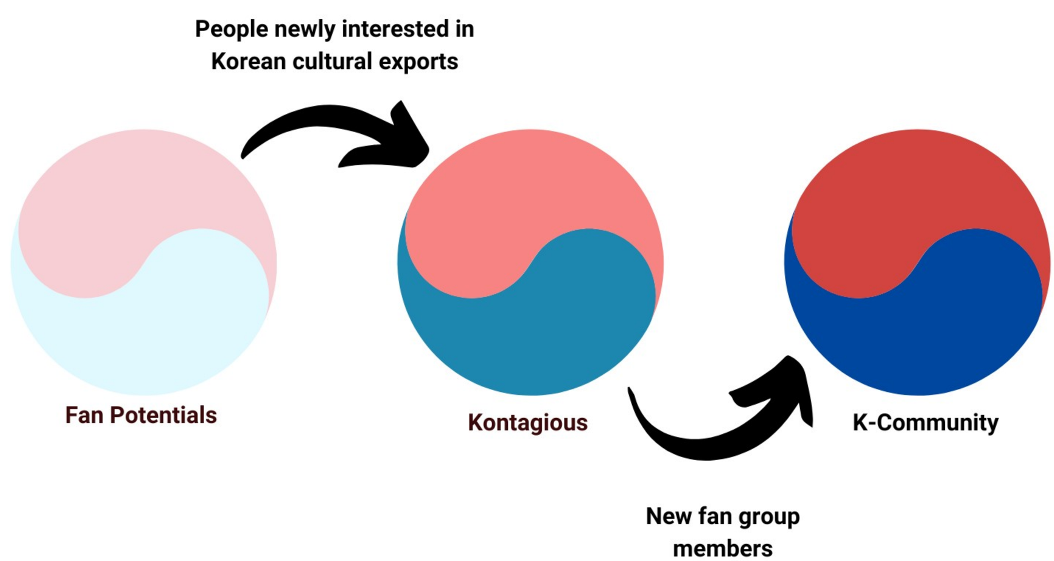

Figure 4.3.1: Visualisation of the SIR Process Applied to the Korean Wave.

To get an insight into the number of people becoming interested in K-content we'll use data from *Google Trends* via the plugin *Glimpse* (Glimpse, 2022) which tells us the number of times a term has been searched on *Google* worldwide. We can assume that each search is done by someone in the **Kontagious** group. Let's look at the two cultural industries that have caused the biggest surge: K-pop and K-dramas.

K-pop artists, also known as *Idols*, are a long-standing feature of the Korean Wave, having an impact since at least the 2010s. Memorably, *PSY*'s *Gangnam Style*, along with its iconic dance moves, went viral and broke multiple records in 2012 (Kim, 2022a). The story is a bit different for K-dramas. Passion for emotive music scores, melodramatic slaps and actors with invisible pores skyrocketed during the pandemic. Lockdowns encouraged people to discover new forms of entertainment, with many falling headfirst into a blissful K-drama binge.

We can fit the function for the size of the **Kontagious** through time to information about the number of searches for K-related content. We will use data from 2010 for K-pop and 2020 for K-dramas. We can minimise the squared distance between the model function and our data, which punishes large deviations more severely, to get the best fit. By finding the line that best suits the data we find estimates for β and γ. Figures 4.3.2 and 4.3.3 show the function fitted to the number of searches for the terms 'Kdrama' and 'Kpop', respectively, on *Google*.

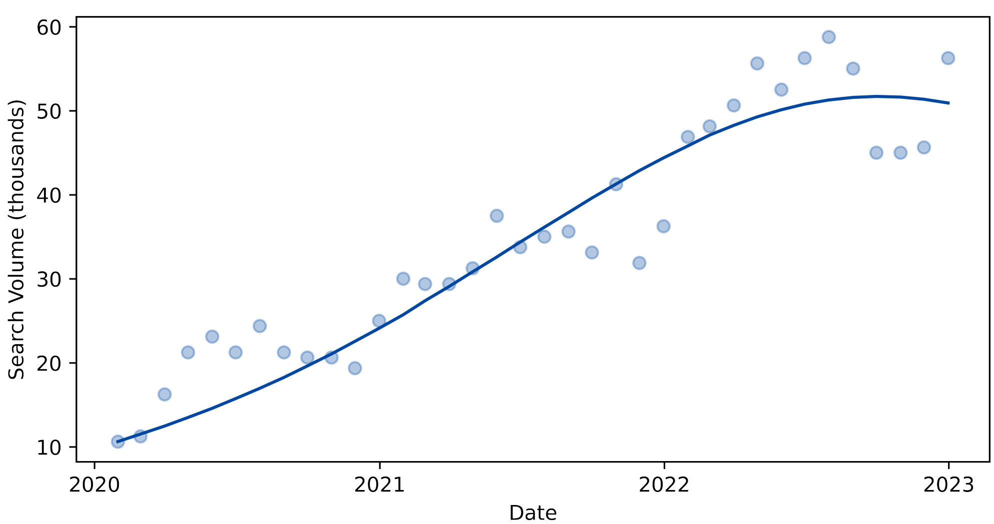

Figure 4.3.2: SIR Model Fitted to the Number of Searches for 'Kdrama' on *Google* from 2020.

There are a number of different search engines as well as streaming, video and social media platforms people can use to search for content. Interested individuals may also

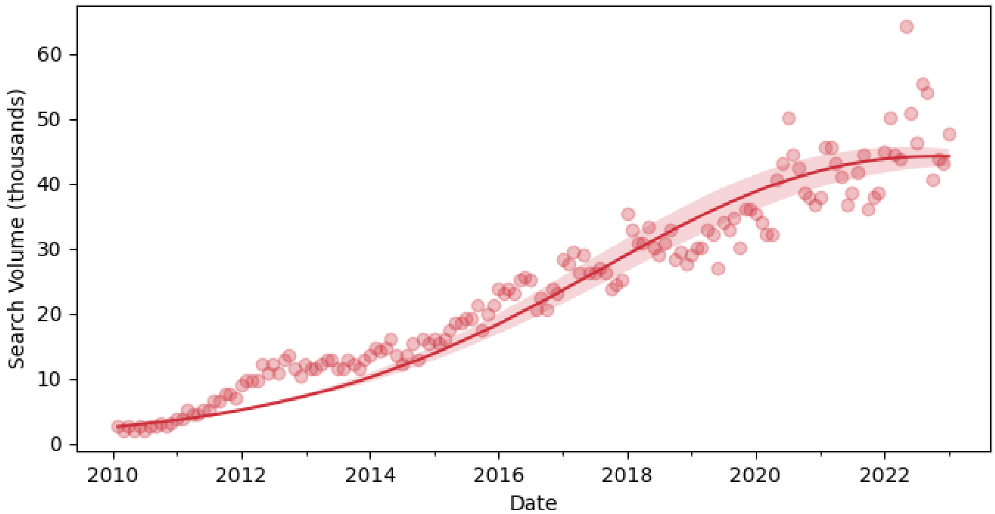

Figure 4.3.3: SIR Model Fitted to the Number of Searches for 'Kpop' on *Google* with 90% Confidence Interval Around β.

search specifically for artist, song, or drama names. As the data we have is only a part of the picture, we compare lots of different estimates for the population size and pick the one that gives us the lowest total error when fitting the model to our data. We assume that there are no members of our population in the **K-community** at the start of the process and that the number of **Kontagious** people at that point is equal to the earliest search volume we have.

It is worth mentioning that more complex models could be used to obtain more accurate estimations, by relaxing some of our assumptions. For instance, taking into consideration that a member of the **K-community** could again become a **Fan Potential** - after all, a K-pop resister could still be tempted to try some mouth-watering tteokbokki! In a stochastic version of this model, the number of newly **Kontagious** individuals could be represented as a random variable. Interested readers might find books on modelling infectious processes and the *shinySIR* package in R useful resources for further investigation. *Significance Magazine* has also published articles offering a more technical breakdown of the SIR model, which make for a fantastic read!

The K Number

The expected number of people who become interested in K-dramas or K-pop as a result of just one fan in the **Kontagious** group tells us important information about the wave (Zhuhadar, 2020). The K Number (known in epidemic modelling as the R number), is the expected number of interactions with **Fan potentials** just one **Kontagious** person has before they become a part of the **K-community**. In our analysis, we calculate this by multiplying β by the average amount of time someone stays in the **Kontagious** group, $1/\gamma$. If our K Number is greater than 1 then it is likely that the number of **Kontagious** people will explode, as each member of the **Kontagious** group is sharing their love for K-content with more than one other person on average.

We can form confidence intervals for β and γ such that we are 90% confident that all plausible values of the true parameter lie between the upper and lower bounds of the interval. Using these we calculate 90% confidence intervals for the K number. For K-pop the K number lies between 3.37 and 4.21. For K-dramas, it is greater than 1.10 but less than 1.79. It seems that K-pop is more infectious than K-dramas. The number of people just one obsessed fan influences is 3.76, on average. For both instances, the lower bound is greater than 1. This tells us that interest has spread and that the passionate K-wave fan base has grown rapidly.

Hallyu Herd Immunity

The height of the Korean Wave will occur when the total proportion of people in the **Kontagious** and **K-community** groups exceeds a threshold. This is when we are at the peak of the K obsession and the number of new people interested in Korean culture exports will start to decline. This is known as achieving herd immunity. The proportion of the population that needs to join the Hallyu herd for this to be reached is $1 - 1/K$ (Zhuhadar, 2020).

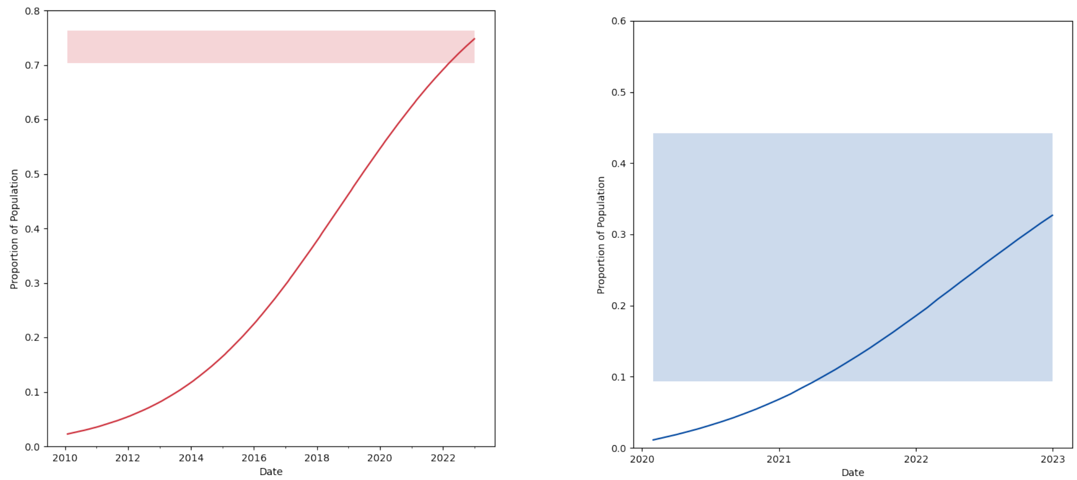

(a) K-pop (b) K-dramas

Figure 4.3.4: Proportion of the Population in the Kontagious and K-community Groups and 90% Confidence Interval Around the Herd Immunity Threshold (Shaded Area)

For K-pop, the region for which Hallyu herd immunity is achieved starts at 0.70 and ends at 0.76. This means that with 90% confidence the total proportion of the population lucky enough to have experienced K-pop bops needs to be between 70% and 76% before we see a decrease in the number of new fans. For K-dramas, the threshold is much lower, between 9% and 44%. Figure 4.3.4 shows that the proportion of the population that has experienced both aspects of K-culture according to our model is within the region for the herd immunity threshold. It looks as if, if we haven't reached it already, the peak of the Korean Wave may be on the horizon.

Surf's Up

Many have marvelled at the sharp rise in popularity of Korean content within the last couple of decades. K-pop, K-dramas and so many other Korean cultural industries have deservedly amassed millions of loyal fans from across the globe. The buzz around K-content, which has led to sharing both online and in the real world, has fuelled the rapid increase of Korean Wave supporters. With the height of passion for Korean cultural exports only just now in touching distance, the influence of the Hallyu fandom will be present for some time. Statistics is wonderfully placed to help us make sense of this cultural phenomenon. After all, the relevance of numbers here cannot be understated. As Kim Young Sam knows better than the rest of us, sometimes even the simplest statistic has the power to make waves.

4.3.3 Article: Could Mathematics Be Big at the Box Office?

This article explores the topic of forecasting to predict box office figures. The application is grounded in the real-world context of the impact of COVID-19 on business and is motivated by the steady re-opening of cinema venues following UK lockdowns. The article was written for the audience of *Mathematics Today* (produced by the *Institute of Mathematics and its Applications*) which is comprised of mathematical graduates. The article is presented below.

Could Mathematics Be Big at the Box Office?

Since the 1890s, storytelling through moving images has gripped audiences across the globe. In post-war Britain, desperate to feel the magic of the big screen once again, people flocked to the cinema in record-breaking numbers. In 1946, 1.6 billion admissions were recorded in the UK (Thompson et al., 2012). In a similar way, after nearly two years of disruption, cinemas expect to see admissions bounce back post-pandemic from their lowest numbers in at least eight decades. This is not surprising. Ultimately, whether it be by bracing for a jump scare, laugh-crying into our popcorn, or swooning for actors that meet impossible Hollywood beauty standards, a visit to the pictures is a way in which we enjoy ourselves and let loose. The stories we are captivated by provide an entertaining distraction and a powerful form of escapism.

Cinema is also a universal communicative tool. The stories told on the big screen comment on and challenge societal issues. Big-screen storytelling has the power to inspire real change within our communities. After all, cinematic experiences are at the heart of the communal events we may have taken for granted before the pandemic. We share feelings, like our suspicion of *Norman Bates*, the frustration at how the door is large enough to float *Jack* as well as *Rose* and the bittersweetness of an adorable alien finally 'phoning home'. Now in the wake of the pandemic, getting the best movies back into our cinemas, so we can once again share in the moments of wonder, shock and awe, feels more important than ever.

Although we know mathematics is a powerful tool, art and mathematics can seem unlikely partners. Among the things for which we do not yet have an exact formula, is what makes a successful and beloved flick. What will satisfy audiences and, of course, what will make money are the most pressing questions for a mathematician interested in the art of the big screen. When trying to understand the complex relationship between audiences and what they love, forecasting is a technique a mathematician has up their sleeve. In the following, we will be taking a quick tour through the forecasting models that can be applied to box office revenue prediction. We will be looking at the mathematical techniques that could help satisfy a deprived post-pandemic audience. In doing so, these methods aim to answer crucial questions: do we have the next Blockbuster? Or just another box office flop?

First to premiere, is *Multi-Linear Regression* which is a fundamental forecasting algorithm, falling into what are known as statistical learning methods. It works by fitting a function between a dependent target variable and independent input variables which describe characteristics. In this context, the target variable y is the box office earnings of a new motion picture. Input variables communicate acquired information about the new movie, such as the number of positive reviews, the value of the budget, box office revenue from the opening weekend and the number of mentions on a social media platform. For n input variables, the box office earnings for a motion picture is given by

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + ... + \beta_n x_n$$

where the β values are regression coefficients. An example of *multi-linear regression* is given in Figure 4.3.5 using data on budget, online popularity and box office revenue.

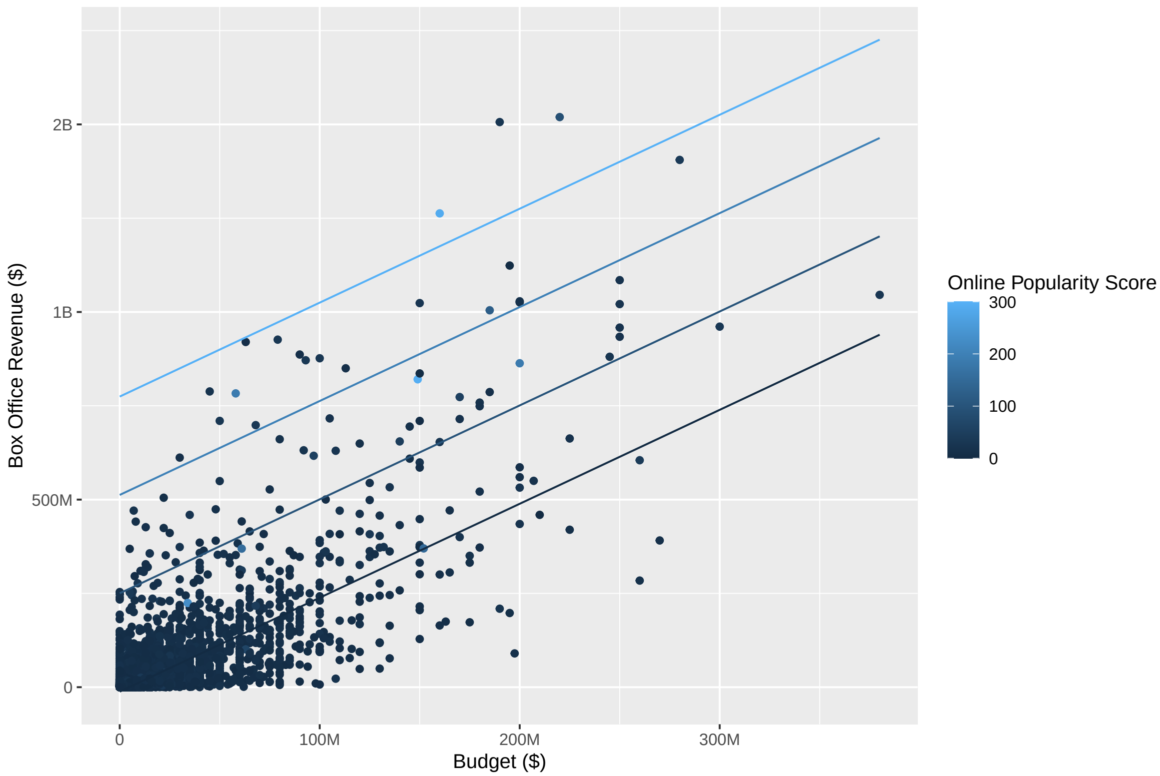

Figure 4.3.5: Multi-linear Regression Example using Budget and Online Popularity Obtained using Data from The Movie Database (The Movie Database, 2021)

To obtain the regression coefficients we can minimise the squared error between actual and predicted box office revenues (Kim et al., 2015). *Multi-Linear regression* can consider a wide range of influences on box office takings but it can only capture linearity in the relationship between the independent features and the dependent target value.

An alternative is to look at a type of time series forecasting models known as new product diffusion models. The *Bass Diffusion Model* (Bass, 1969) is especially popular and is not restricted by linearity. It considers two factors, innovation and imitation.

The innovation parameter represents all mass media communication of the film. The imitation parameter explains responses of audiences that have already seen the movie and portrays the effect of word of mouth (Dellarocas et al., 2007). Figure 4.3.6 provides an illustration. The model gives a hazard rate which is the probability that a potential audience member that has not yet seen the film will watch it at time t. The hazard rate is given as:

$$h(t) = \frac{N'(t)}{1 - N(t)} = p + qN(t)$$

where p is our innovation coefficient and q is the imitation coefficient. $N(t)$ is the cumulative portion of potential moviegoers that have seen the film at time t. That is, the number of people that have seen the film by time t out of all possible cinemagoers, P. $N'(t)$ is its derivative. What the hazard rate tells us is that the portion of the potential moviegoers that go to see the film for the first time at time t can be written as a linear function of the people that have already seen it. The *Bass Diffusion Model* differential equation is

$$N'(t) = p + (q - p)N(t) - q\left[N(t)\right]^2 .$$

At time $t = 0$, the cumulative proportion of the potential audience that has seen the film is zero. We can solve the differential equation and write $N(t)$ as

$$N(t) = \frac{1 - e^{-(p+q)t}}{1 + (q/p)e^{-(p+q)t}}.$$

The model has three unknown parameters which need to be estimated: the innovation parameter, p, the imitation parameter, q, and the number of total potential cinemagoers, P. The latter, P, can be estimated in a number of ways using historical data such as total cinema admissions. When we look at the number of new moviegoers, we are effectively splitting the audience of a film into two categories: innovators and imitators.

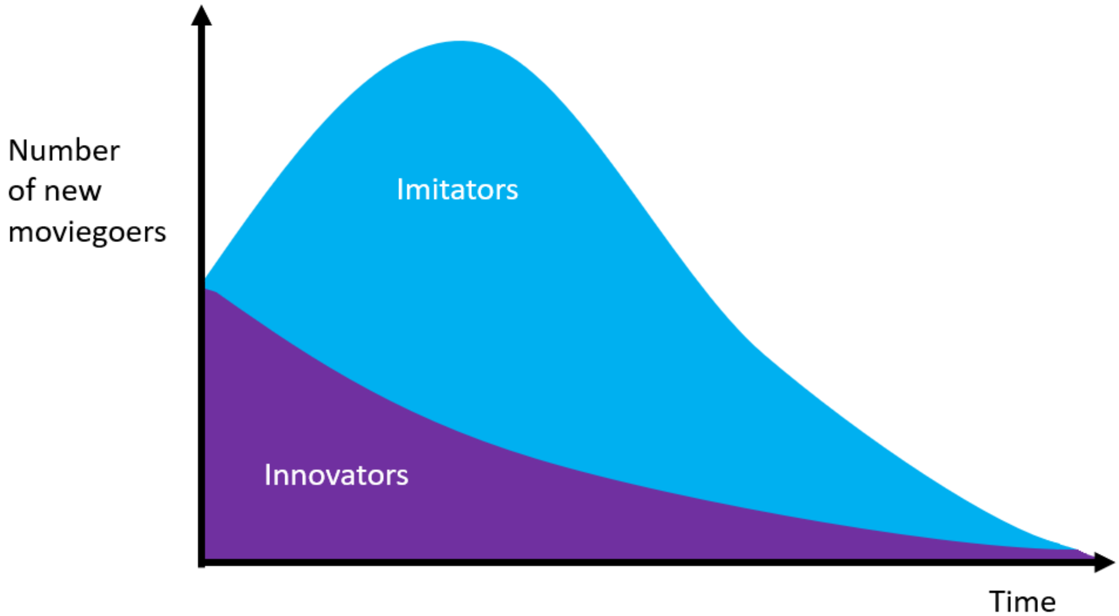

Figure 4.3.6: Bass Diffusion Model: Illustration of New Moviegoers

Innovation is what gives a movie its momentum from the get-go. These are the factors that entice audience members from early on, whether that be studio reputation, advertising and promotions or just the presence of *Tom Cruise*. In practice, the box office takings from the opening weekend can usually be used to gain a reasonable estimate of p. Conversely, the imitation parameter needs to capture all the internal factors likely to get you off your couch and into the cinema, chomping on popcorn. This includes word of mouth as well as social media buzz. In reality, one possible way we could approximate this is by measuring the positivity of user reviews online.

Once we have $N(t)$ we can multiply by P to get the cumulative number of people who will watch the film by time t. Using information about ticket revenue we can calculate the predicted total box office figure. In Figure 4.3.7, a *Bass Diffusion model* dispels the mystery around revenue for the whodunnit *Knives Out* (Box Office Mojo, 2019). Admissions for the blockbuster were halted on its 106th day, at the beginning of the pandemic.

Another option to model the expected revenue for a new motion picture is to venture

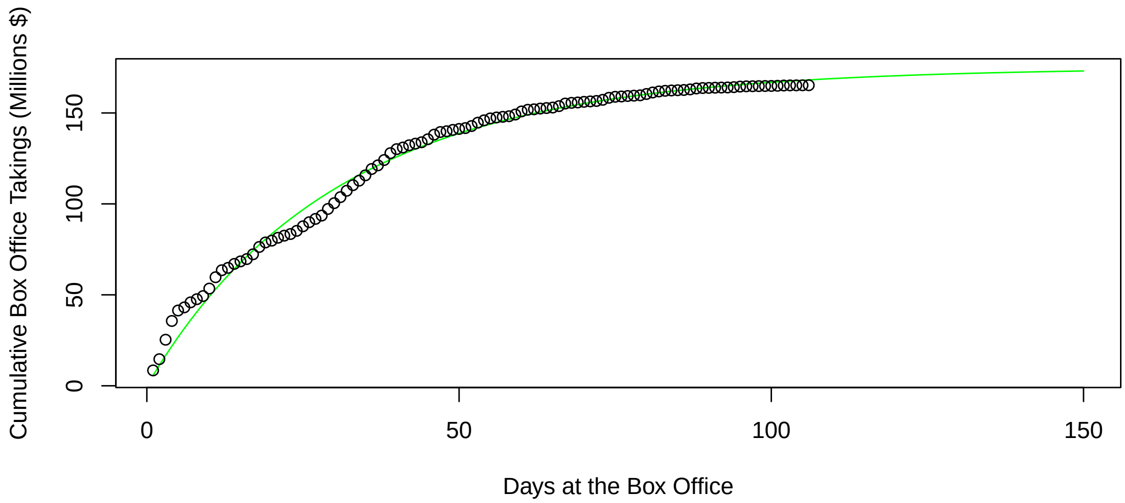

Figure 4.3.7: *Bass Diffusion model* Fitted to Box Office Revenue Data for *Knives Out* (Box Office Mojo, 2019). Predictions are Given by the Green Line.

into machine learning-based regression algorithms, such as the *k-NN* algorithm. The *k-NN* process uses information about previous motion pictures and their corresponding box office takings. By looking at previous movies with similar features, *k-NN* utilises known box office earnings to forecast the success of a new release.

The first step is to calculate how similar the new release is to previous films. For instance, we could measure similarity based on characteristics such as budget and rating. One way to measure the distance between a film and a new release is to compute the Euclidean distance, the square root of the sum of the squared difference between the new point and the existing points (after suitable normalisation). The set of the k most similar films are selected and dubbed the 'nearest neighbours'. For each of the nearest neighbours, the box office takings, y_j, are observed and weights, w_j, are assigned.

The predicted box office earnings for a new release i are given as the following (Kim et al., 2015):

$$\hat{y}_i = \sum_{j \in \text{k-NN}(\mathbf{x}_i)} w_j y_j.$$

To execute *k-NN*, the number of neighbours and the weights assigned to those neigh-

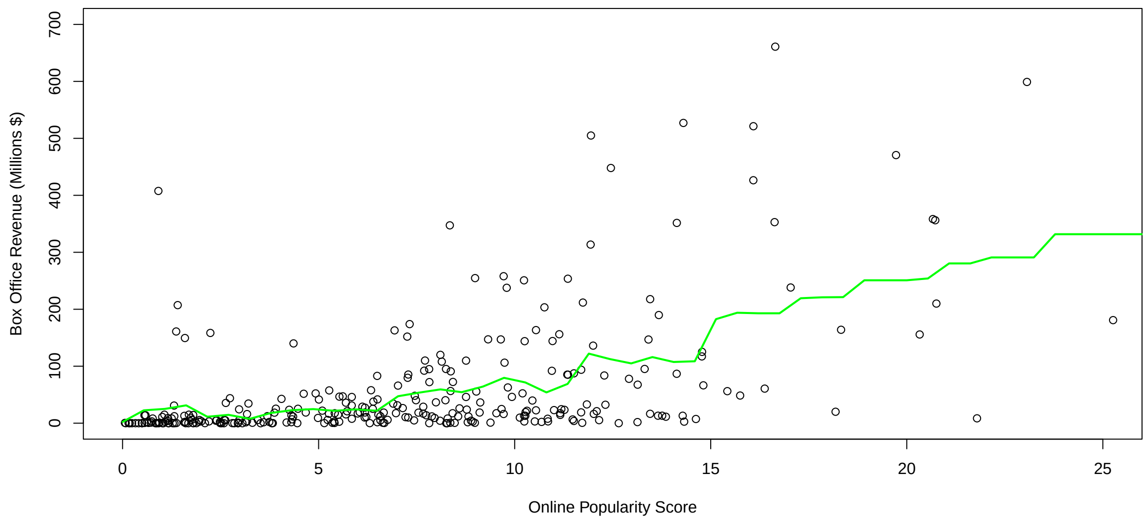

Figure 4.3.8: Example of *k-NN* Using 300 Samples Obtained from The Movie Database (2021) with $k = 20$. Predictions are Depicted by the Green Line.

bours need to be chosen. The number of neighbours, k, is often chosen by evaluating the performance of the method for multiple different k values. When it comes to the weights, it feels logical to assign a neighbour that is further away a smaller weight than a neighbour that is closer. This commonly leads to the adoption of some kind of kernel function, which monotonically decreases as distance increases. Figure 4.3.8 shows box office predictions using data about online popularity for 300 movies from The Movie Database (2021) and a value of k equal to 20.

And that's a wrap! Predicting the success of art, and how to satisfy worldwide audiences, is an inherently complicated task. By using forecasting algorithms, mathematics and storytelling can come together with the aim of maximising audience gratification. Soon cinemas will get back on their feet and it will feel safe to once again journey into alternate universes and join breathtaking adventures. We can only hope that box office hits are all that lie in store for the patient audiences that have longed for their return to the big screen.

4.3.4 Talk: Back to the Future with Mathematical Techniques

This talk recounted the author's journey in mathematics education and shared mathematical techniques used in forecasting and prediction. The mathematical techniques introduced are outside of the standard school curriculum and have an emphasis on solving problems with interesting applications. Within the design, the author incorporated themes aimed to challenge preconceptions that deter students from pursuing mathematical career pathways. The length of a standard delivery of the talk is 45 minutes. However, reflective of the flexible demands of outreach events, the delivery of the talk has also met alternative timing requirements.

To begin, the author introduced participants to the subject of OR. The author provided a high-level overview of the work they have done in their PhD and mentioned the real-life applications of their research. Following this, the author described their journey to studying for a PhD, in doing so highlighting themes related to ability and learning. In discussing the experience of learning mathematics at school, the author shared that they found mathematics both enjoyable and challenging. Relative to the inaccurate perception of mathematicians as individuals who find mathematics easy, the author relayed formative advice from university lecturers related to the unavoidable challenge of learning mathematics. An example can be seen in Figure 4.3.9.

The author introduced the topic of their final year undergraduate project, using Game Theory to model and predict the outcome of narratives within storytelling. The author described the concepts of a two-person strategic game, including the sets of players, actions and player preferences (Osborne, 2004). The author introduced a proposition within Game Theory, the Theory of Moves (Brams, 1993). The author explained that by the Theory of Moves, a game begins in an initial state and ends in a nonmyopic equilibrium, defined as a state from which neither player would have reason to move and identified using players' preferences. During the game, a player can change the state of the game from the initial state to a more favourable state by unilaterally

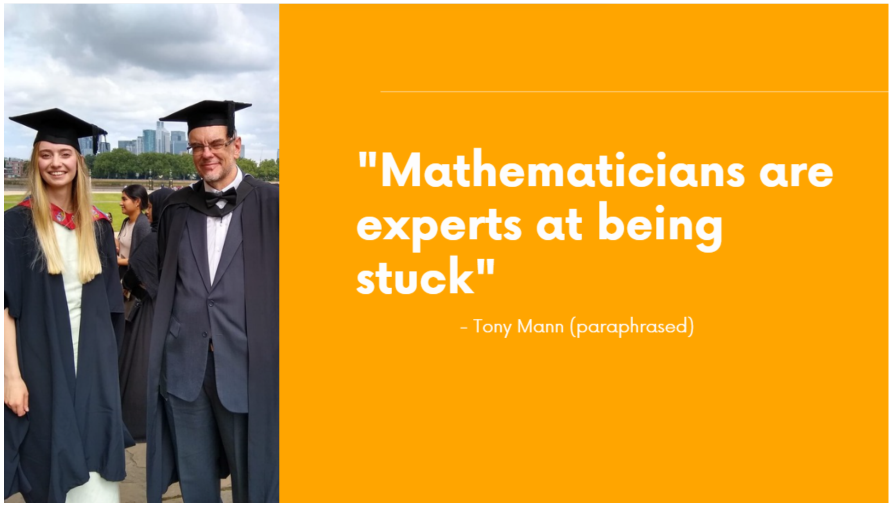

Figure 4.3.9: Slide from 'Back to the Future with Mathematical Techniques' Talk.

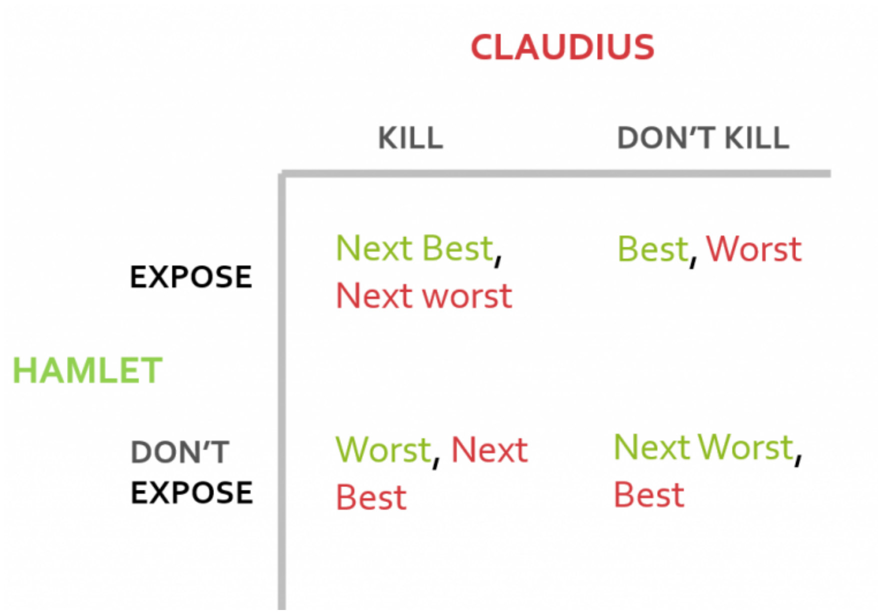

Figure 4.3.10: Game Theoretic Interpretation of Conflict within *Hamlet* (Brams, 2011) as Shown in 'Back to the Future with Mathematical Techniques' Talk.

switching their action.

Once core concepts were understood by the audience, the author demonstrated how the Theory of Moves can be applied to the central conflict between the protagonist and the villain within *Shakespeare*'s *Hamlet* by depicting it as a two-person game, as shown in Figure 4.3.10 (Brams, 2011). The author communicated the following details. The players of the game are *Hamlet* and *Claudius*. *Hamlet*'s set of actions includes choosing whether or not to expose *Claudius* as the King's murderer. *Claudius*'s action are to kill or not kill *Hamlet*. The states of the game are preferentially ranked from best to worst for each character. The states are summarised below.

1. *Hamlet* exposes *Claudius* as the King's murderer, *Claudius* kills *Hamlet*
2. *Hamlet* exposes *Claudius* as the King's murderer, *Claudius* does not kill *Hamlet*
3. *Hamlet* does not expose *Claudius* as the King's murderer, *Claudius* kills *Hamlet*
4. *Hamlet* does not expose *Claudius* as the King's murderer, *Claudius* does not kill *Hamlet*

The author informed the audience that the game begins in the fourth state and that by the rules of the Theory of Moves, *Hamlet* switches action, changing the state of the game to the second state. In response, *Claudius* changes the state of the game to the first state. In this state, neither player can benefit from changing the state of the game and thus the demonstration leads to a nonmyopic equilibrium, signalling the end of the game. The outcome corresponds to the true ending of the play in which *Hamlet* exposes *Claudius* as a murderer but ultimately succumbs to a wound caused by a blade poisoned by *Claudius* (Brams, 2011). Thus, the mathematical analysis provided an accurate prediction of the resolution of the conflict.

Following this, the author relayed an additional example from the popular television show *Game of Thrones* and characters *Cersei Lannister* and *Daenerys Targaryen*. The conflict central to the final season of the show involves these characters' decision to form an alliance. Figure 4.3.11 depicts the conflict as presented in the talk.

In relation to the third design principle, relatability, the author discussed feeling anxiety about exploring applications of Game Theory in storytelling, given that the application aligned with their non-scientific interests. However, the author shared a positive response from their supervisor. Further, the author mentioned the reassuring

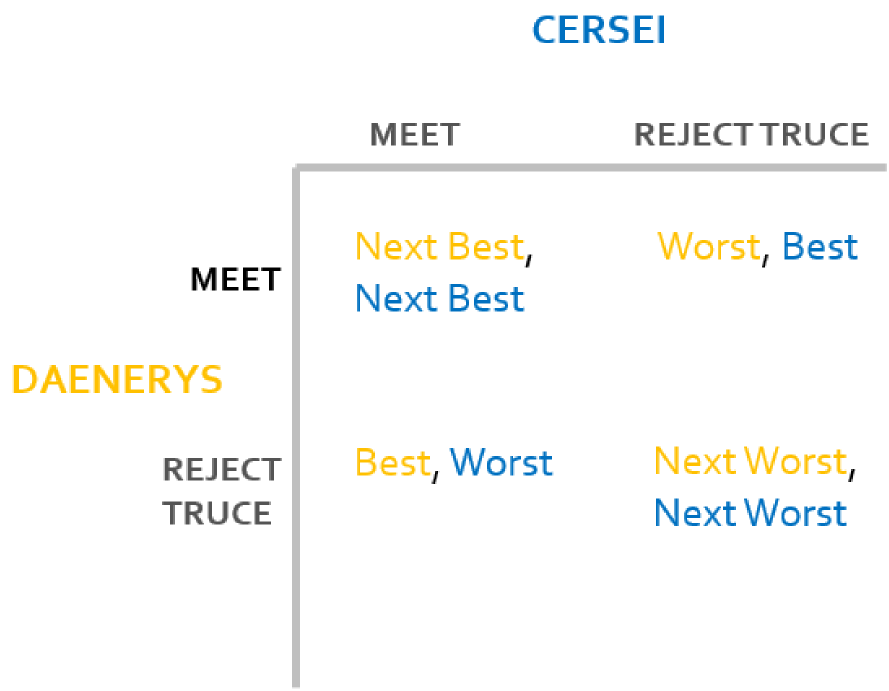

Figure 4.3.11: Game Theoretic Interpretation of Conflict within *Game of Thrones* as Shown in 'Back to the Future with Mathematical Techniques' Talk.

reception from other mathematical teaching staff, the high assessment outcome and their enjoyment of completing the project.

Near the end of the talk, the author highlighted the exciting applications of mathematics. In discussing how mathematical techniques can be used to predict the future the author gave advice designed to inspire the students. The author communicated that in the future the students are also likely to discover surprising applications of mathematics and oftentimes find learning new things difficult. The author reminded the students that feeling stuck when trying to understand something new in mathematics is a common feeling and that learning mathematics is inevitably challenging. The author emphasised that, for many mathematicians, the difficult nature of building an understanding of mathematical topics is in large part what makes learning enjoyable. The talk finished with the opportunity for the audience to ask questions.

4.3.5 Workshop: Mathematics at the Box Office

A workshop was designed for the Royal Institution (Royal Instituition, 2024) masterclass series. Participants were of GCSE level and elected by their school to attend events. The workshop was based on exploring forecasting techniques that can be applied to predict the monetary success of films at the box office. The workshop was 120 minutes long.

To begin, the author introduced themselves as a mathematician and PhD student. Towards the first design principle, engagement, the author engaged the audience by asking if they knew what a PhD student is. Following audience suggestions, the author discussed the PhD qualification and, more widely, the aims of mathematical research. An overview of OR was also given. The author went on to describe their PhD research topic at a high level, focusing on the application to real-life impact in reducing waste and making sustainable decisions.

The author set the motivation for the problem that was explored in the workshop by discussing the negative impact of the COVID-19 pandemic on cinemas, citing the bankruptcy of the US chain of *Cineworld*. The author explained that, once deemed safe to reopen, cinemas struggled to account for the losses incurred from extended closures. This topic was selected because it was a real-world, high-stake problem relevant to the interests of the general population. The author proposed that forecasting techniques might be able to help cinemas determine the popularity of a film before its release, thus allowing cinemas to tailor their screenings and secure profit. The author relayed that this analysis could also be useful for stakeholders to gauge whether a film is a worthy investment before production begins.

The author delivered an introduction to Linear Regression and Least Squares Optimisation. With a focus on the second design principle, communication, the techniques were described within the context of the problem setting, by using the budget of a film as an explanatory variable and the box office takings as the response variable. The author used a graphical representation to build audience understanding.

The students were then given a small sample of real data on the budget and US box office takings of eight films from 2022. The students were tasked to model the relationship between budget and US box office takings. Using their model, they were asked to make a prediction for how much a new film, *Avatar 2: Way of Water*, would make at the box office, given its particularly large budget of $460 million US dollars. The data provided to students is presented in Table 4.3.1. Figure 4.3.12 shows a related slide from the talk.

Table 4.3.1: Film Budget and US Box Office Takings Data Provided to Students in Workshop

Film	Budget (Millions of US Dollars)	US Box Office Takings (Millions of US Dollars)
Sonic the Hedgehog 2	110	190
Uncharted	120	148
Elvis	85	151
Morbius	75	73
Puss in Boots: The Last Wish	25	70
Everything Everywhere All At Once	25	70
DC League Super Pets	90	93
The Woman King	50	67

Later, the students were introduced to more sophisticated forecasting models, including the Bass model (Bass, 1969), which they implemented to find updated predictions. The author asked the students to interpret the parameters, provoking discussion about the relationship between film characteristics and its box office performance. In

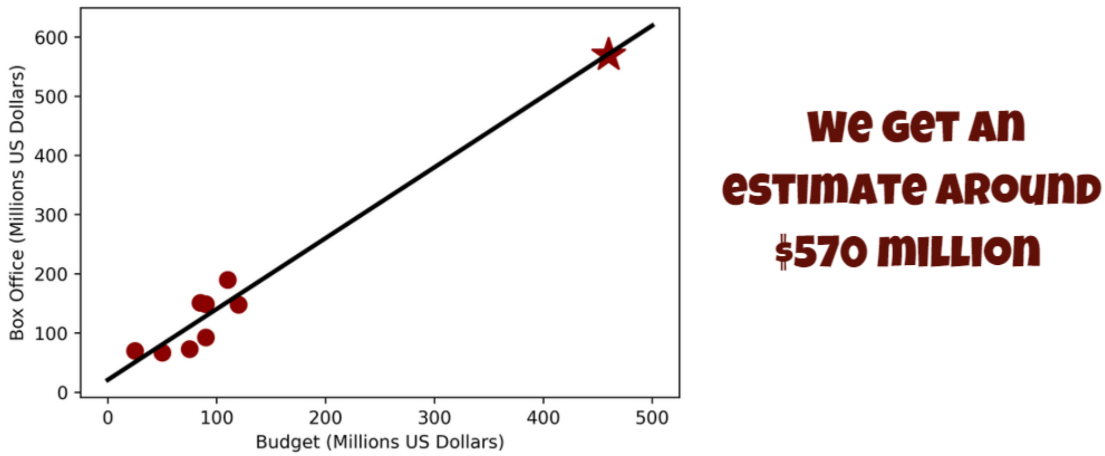

Figure 4.3.12: Slide from 'Mathematics at the Box Office' Workshop

this way, students were encouraged to link their findings to the real world and explore the interdisciplinary nature of the problem. The author asked the students to consider the limitations of the methodology. Discussion prompts included the differences between linear and non-linear models and ways to improve predictions (i.e., increasing the amount of data used or including additional explanatory variables).

The workshop finished with a 20-minute question-and-answer session. The student's questions were varied. Most questions related to the experience of studying mathematics at university and possible career pathways. Some students wanted to know more about the author's personal career experiences and obstacles they had faced. In particular, the author was asked for advice on dealing with feelings of self-doubt or low confidence whilst studying mathematics.

4.4 Reflection

This section offers a discussion about the impact of the work presented in this chapter. Firstly, a record of STEM outreach activity completed by the author from December 2021 to June 2024 is provided. Thereafter, feedback from teachers on the delivery of the outreach activities is reviewed. The experiences of the author are used in combination with comments made in feedback to analyse the effectiveness of this dissemination in achieving the aims of outreach. Additionally, a reflection on the inclusion of content designed to encourage diversity within the STEM community is discussed.

4.4.1 Record of Activity

The in-person activities conducted by the author vary in location, style, length, audience demographic and audience size. In recognition of the disparity in the fields of mathematics, statistics and OR, some events were tailored to encourage students belonging to particular demographics, such as the Enrichment Day run by the Advanced Maths Support Programme which invites attendees of underrepresented genders. The majority of the activity took place in the North West of England with the target audience being students studying for A-Levels. A talk was the most common style of delivery. Further details of the audiences and location of the in-person STEM outreach events are outlined in Table 4.4.2.

A record of published work is given in Table 4.4.1. The article given in Section 4.3.2 was published in *Significance Magazine* which aims at a general audience with differing levels of mathematical and statistical knowledge. A second article, presented in Section 4.3.3, was published in *Mathematics Today*, the magazine produced by the *Institute of Mathematics and its Applications* for readers who are graduates of the mathematical sciences.

Table 4.4.1: Record of Published Activity

Title	Publication	Date Published	Audience Demographic
Surfing the Korean Wave	*Significance Magazine*	Dec-23	General Audience
Could Mathematics be Big at the Box Office?	*Mathematics Today*	Dec-21	Mathematics Graduates

A description of the in-person outreach events are presented below.

Greenwich Maths Time

Greenwich Maths Time is a day of talks, workshops and interactive activities held at the University of Greenwich in London designed to allow Year 9 and 10 students to explore mathematics. In 2024, it was run as a part of the *Institute of Mathematics and its Applications* Festival of Mathematics. The event aims to allow students to explore maths used in everyday life, to encourage students to consider career pathways in mathematics and to showcase the diversity of the mathematical community (University of Greenwich, 2023).

Lancaster University Year 12 STEM Taster Day

The STEM Taster Day is a school outreach initiative run by Lancaster University offering Year 12 students the opportunity to discover more about studying STEM-related subjects at university. The event is cross-departmental, with hands-on sessions offered by enthusiasts and communicators from many STEM disciplines. The day aims to inform students about the impact of STEM research with an emphasis on practicality (Lancaster University, 2023b).

The Advanced Mathematics Support Programme Steps to University Mathematics Enrichment Day

A government-funded initiative the Advanced Mathematics Support Programme

(AMSP) aims to improve participation and teaching in mathematics at core and post-16 levels. Support provided by the AMSP is given to areas of low social mobility and where participation in A-level is low with the aim to encourage and support all students regardless of their location, gender, ethnicity, or background. The Steps to University Mathematics Enrichment Days happen nationally and are designed for female, non-binary and transgender Year 12 students considering studying mathematics at university. These events connect students with early-career women mathematicians and run in various locations across the UK (Mathematics Education Innovation (MEI), 2018).

Royal Institution Masterclass

Royal Institution masterclasses are extracurricular interactive workshops run across the UK, typically for secondary school students on Saturday mornings. The masterclasses are interactive and cover a wide range of STEM subjects within the fields of mathematics and computer science. The masterclasses aim to nurture students' engagement with STEM, allow students to build confidence in STEM skills and excite students by exploring topics not taught typically taught within the classroom. By connecting students with STEM experts and enthusiasts many students are inspired to consider a career within STEM (Royal Instituition, 2024).

Florence Nightingale Day

Organised by Lancaster University, the Florence Nightingale Days are an effort to promote mathematics and statistics by showcasing successful women mathematicians and statisticians at a variety of stages of the career ladder. Aiming to inspire, the next generation of mathematical scientists, the event emphasises the variety of breadth of possibilities pursuing a career in mathematics can bring. The event is named after Florence Nightingale, who is most famous for her work as a nurse, in recognition of her lesser-celebrated statistical work. She pioneered data visualisation techniques that saved countless lives. (Lancaster University, 2023a).

Table 4.4.2: Record of In-Person Activity

Event	Date	Activity	Length (minutes)	Audience Size	Audience Demographic
Greenwich Maths Time	Jun-24	Talk	50	300	GCSE & A-level students
Lancaster University STEM Taster Day	Jun-24	Talk	45	100	A-level students
Greenwich Maths Time	Jul-23	Talk	50	300	GCSE & A-level students
Lancaster University STEM Taster Day	Jun-23	Talk	45	100	A-level students
AMSP Steps to University Mathematics Enrichment Day (London)	Apr-23	Talk	25	60	A-level students
AMSP Steps to University Mathematics Enrichment Day (Lancaster)	May-23	Talk	25	60	A-level students
Royal Institution Masterclass	Jan-23	Workshop	120	60	GCSE students
Florence Nightingale Day	Jan-23	Talk	45	300	GCSE & A-level students

4.4.2 Feedback & Author Experiences

Teacher feedback was obtained from a subset of the outreach activities completed by the author. The comments are labelled by letters A-H and presented in Table 4.4.3. With the exception of comment E, which was communicated directly to the author, the feedback presented in Table 4.4.3 was collected by event organisers and passed on to the author anonymously.

Table 4.4.3: Teacher Feedback on Author's Outreach Activity

Label	Event	Comment
A	Florence Nightingale Day	'I thought the talks were great. I'd love to have the second speaker (Robyn) come and give a talk to students at our school!'
B	Florence Nightingale Day	'All the students really enjoyed the talk from Robyn especially regarding her pathway into maths'
C	Florence Nightingale Day	'Robyn's talk - they thought she was interesting and really engaged with what she was talking about - felt they could relate to her.'
D	Florence Nightingale Day	'The girls were so inspired by the women speakers, it was a joy to listen to how excited they were on the journey home. One girl now wants to get into mathematical modelling following this session. They really were inspired to hear how mathematics could be used in so many fields. We had a great day and would love to bring a group every year!'

E	Florence Nightingale Day	'I came to the Florence Nightingale day a couple of years ago and heard you speak, and I have always remembered your talk and how much my students were fascinated by it. I have a new group of Year 12 students, who are all very interested in maths and have been asked to look at an inquiry project based on a question they found interesting. I was telling them about your talk on that day and now they are going away and looking into game theory and prediction modelling. I have shown them some of your articles and they are fascinated so I just wanted to pop an email across to say thank you for helping me to make them more interested in a maths inquiry and developing their skills.'
F	Greenwich Maths Time	'All the speakers were great and the students were inspired especially by Robyn. I think it made them feel at ease knowing that studying Maths is tough and you will get stuck but you need to keep going.'
G	Greenwich Maths Time	'The 'Back to the Future' talk was especially impactful for our students as we are a vocational film school but we have a small a group who are studying A-Level Maths so it really resonated with them. Two of them were so inspired by the talk that they made the decision to stick with the A-Level course next year when they were considering dropping.'

H	Lancaster University STEM Taster Day	'She (A participant) came out absolutely buzzing about the maths session in particular, which is particularly impressive because it was her 2nd choice. It sounds like they have a couple of great activities, which were really thought-provoking. If you get chance, let them know that it changed at least one attendee's perspective on maths!'

The comments in Table 4.4.3 highlight that the work reported in this chapter has been well received and contributed to the aims of promoting interest in mathematical subjects. Comment A indicates that the outreach activity was beneficial for students. The excitement held by the participant mentioned in Comment H insinuates that the sessions delivered by the author were enjoyable and that activities based on forecasting and prediction inspired a positive view of mathematics, statistics and OR. Accordingly, the author found that students engaged with content, expressing a willingness to understand mathematical concepts. Following sessions, students waited to speak one-on-one or in small groups with the author, demonstrating a keenness towards mathematical topics. Comment G reflects a particularly desirable outcome in which students were inspired to pursue a pathway in mathematics after previously deciding not to. This provides a clear argument for the effectiveness of this work towards the long-term goal of improving student retention.

The feedback and the author experiences also reflect that the author's activity allowed students to build competency in mathematical skills. The mention of mathematical modelling in Comment D showcases that the activity allowed students to develop an understanding of the breadth of mathematical fields. A progression in ability is also insinuated by Comment E, with specific reference to the building of skills about both in-person activity and written communication completed by the author. At the end of many in-person events, students sought recommendations for additional learning resources. This indicates that this work has had a positive influence on the development of mathematical skills.

The deliberate incorporation of themes relating to mathematical belonging in this work has highlighted that irrespective of enthusiasm for the subject, students are affected by barriers to mathematical belonging. Discouragement from pursuing mathematics, in some cases, originated from the students' beliefs about competence. In discussions with students, many cited an average performance in classroom assessments or experiences of finding learning challenging as reasons to not pursue a career in mathematics. In the delivery of the workshop described in Section 4.3.5, students reacted negatively to making mistakes during the exercises. Comment F reflects the attempts of the author to dismantle inaccurate expectations of ability associated with mathematical professionals. Comment B mentions the inclusion of the author's personal journey, highlighting the importance of communicators to retell experiences of setbacks as well as successes. This aspect of the activity had an impact, emphasising the importance of outreach initiatives to diversify perceptions related to the learning of mathematics.

In the author's experiences of delivery, imposing transparency about belonging and diversity within mathematics revealed insights into students' perceptions of the mathematical community. In particular, gender stereotypes were identified by the author as a key theme that deterred students from pursuing a mathematical career pathway. Events targeted at students of underrepresented genders (e.g., the *AMSP Steps to University Mathematics Enrichment Day*) evoked open discussion about equity within the classroom. Participants spoke of their hesitancy to answer questions in class in fear of being incorrect and expressed pressure to oppose the stereotype of women not being good at mathematics. These students felt this opinion was endorsed by some of their male peers and shared instances of sexist comments made in the classroom. Comment D shows that girls were positively affected by the author's talk, referring to women speakers having an impact. Comment E evidences the favourable reception of the author's articles for subsequent cohorts. Given that the article presented in Section 4.3.2 explores female-dominated interests (e.g., K-pop, K-dramas), this shows that outreach can have a long-lasting effect on diversifying perceptions related to gender within the mathematical sciences.

In verbal feedback that students gave directly to the author, the notion of having non-scientific interests was a repeated theme. Multiple students relayed their passions for other disciplines, such as fashion, history or philosophy, and were inspired by the interplay between mathematical ideas and non-traditional applications. These students spoke of feeling conflicted about maintaining a passion for mathematics whilst also being interested in non-scientific pursuits. Interdisciplinary applications explored throughout the in-person activities were key to challenging the expectation that mathematical professionals hold exclusively scientific interests. Considering Comment H, the context of film and television is identified to have resonated with some students' passions. In interactions, the author was able to emphasise that qualities commonly associated with humanities and artistic subjects, such as creative thinking, are also essential skills in the study of the mathematical sciences. The applications considered in the articles in Sections 4.3.2 and 4.3.3 (television, film and music) also directly challenge the concept of mathematicians having singular interests. Exposing students to mathematical techniques and applications that are not taught within the classroom is effective not just towards the aim of promoting interest in mathematics but also in challenging preconceived ideas about mathematical identity.

Overall, it is evident from the feedback that the inclusion of content intended to remove barriers to mathematical belonging is necessary and valuable. Despite there being less focus on the aim of mathematical belonging within the workshop in Section 4.3.5, in comparison with the talk in Section 4.3.4, several questions were asked during the question and answer session on topics related to competence. This included questions about how to build mathematical confidence. Comment C communicates that students felt that they could relate to the author, improving the accessibility of mathematical careers. Comment E shows how the outreach activity has had a sustained impact. The talk delivered at Florence Nightingale Day has continued to inspire new cohorts of students and the content has been adopted by their teacher. This supports an argument for the careful design of outreach activity which can leave a lasting impression on future generations.

Although useful, it should be acknowledged that analysis based on the feedback given above has limitations. The feedback is from teachers across three events and cannot be assumed to be representative of the opinions held by all students who attended the sessions. In the author's experience, written feedback from students is rarely collected at outreach events. Feedback from teachers whose students attended outreach activities is much more commonly requested by event organisers however, after collection, it is not always communicated directly to speakers. More rigorous qualitative methods should be utilised by organisers and communicators to formally assess the impact of outreach activity. Longitudinal studies are needed to determine the long-term impact of outreach towards uptake in careers within STEM and solidify the proposition that, in addressing barriers to mathematical belonging, outreach forms valuable intervention.

4.5 Conclusion

Expertise in STEM subjects is fundamental to facing societal and economic problems that have serious global consequences. Although imperative, STEM sectors are facing a crisis of shortages across the workforce. Efforts to increase the visibility of careers and foster interest in mathematical topics are needed to improve the uptake of STEM roles across industry, business and government. Additionally, improving diversity is a necessary focus for the STEM community to ensure future prosperity. In terms of gender, the underrepresentation of women within the community is a long-standing issue. Stereotypes are perpetuated by parents, teachers and wider society making misaligned beliefs about STEM subjects and mathematical science professionals difficult to dismantle.

Outreach is pivotal in efforts to overcome the difficulties facing the STEM sector. Talks, workshops and written communication provide practitioners the chance to share excitement for the mathematical sciences and its applications to broad audiences. Communicators foster engagement with scientific topics, relating core ideas with real-world issues affecting the wider population. Outreach also offers the opportunity for professionals to connect with future generations. Activities incorporate content outside of the standard curriculum which excites students and supports the learning of related skills. In turn, outreach plays a crucial role in encouraging the pursuit of science-related career paths.

This chapter has introduced original content that contributes towards the aims of outreach. The content focused on the communication of forecasting and modelling methodology and has included various formats of dissemination including a talk, a workshop and two written pieces of science communication. The work was developed in line with core design principles. Not all outreach initiatives include the aim of addressing the STEM community's shortage issues and improving recruitment and retention of individuals belonging to underrepresented groups. Towards the aim of improving the representation of women in mathematical fields, the work has also aimed to dispute stereotypes that contribute to gender disparity.

On reflection of experiences and feedback, the author's content was positively received and led to a tangible impact. A notable outcome of this work was improved student retention. These results support the hypothesis that topics within forecasting and modelling can promote interest in methodology and applications of statistics and

OR. A strength of the work introduced is its diversity in terms of activity, with original content provided across three mediums. Moreover, the content has considered a variety of audience demographics, including school students, mathematical graduates and the wider public. The author's focus on designing relatable content featured prominently in teacher feedback, showing the potential of outreach activity in overcoming the issues faced by the STEM community. The author's work lays the foundation for future investigation into the effective design of outreach activities.

The author's work could be enhanced in many ways. In the majority of the content presented, the author used applications to film & television to engage audiences. However, a wider variety of applications could also be explored. One approach could be to focus on applications that are highly relevant to the general public (i.e., climate change). Equally, the dissemination of topics from active research areas could excite audiences. The work given in this chapter considers the disparity between men and women in fields within the mathematical sciences. However, investigation into how outreach can help improve the underrepresentation of people with disabilities, ethnic minorities, the LGBTQ+ community and those from disadvantaged socioeconomic backgrounds is an important line for further work. Additionally, the analysis provided is focused on the experience of one individual within a limited time frame. Longitudinal studies into the effect of outreach initiatives designed to improve diversity in mathematical fields could be conducted to better understand the impact of outreach work in this area.

Chapter 5

Conclusions and Further Work

This thesis addressed challenges related to forecasting and modelling for long-term problems. The work presented centred on the gap in forecasting methodology for after-market products with long-standing patterns. This topic was motivated by the strategic decisions faced by an automotive manufacturer. We also explored the dissemination of forecasting and modelling ideas and discussed the use of outreach to ensure methodology has a long-lasting, practical impact.

In closing, we summarise the key findings of each chapter of this thesis, the limitations of our approaches and valuable avenues for future research.

5.1 Forecasting Declining Demand

In Chapter 2, we introduced methodology to support long-term inventory decision-making. Our approach was inspired by the LTB decision, for which a company has to make a final order that is intended to cover all future demand. We presented a stochastic model to forecast demand that exhibits a long-term declining pattern and may have periods of zero demand, consistent with spare parts in the EOL phase of the life cycle. We represented demand as Poisson with a mean that declines geometrically. The model specifies two parameters, an initial mean demand size where the decline in demand begins and a rate of decline. The parameters are estimated using maximum likelihood estimation. Theoretical insights were obtained which expose a disparity in the performance of the model, by which better forecasts are determined for products of higher demand volume. To improve forecasts for slower-moving products we contributed an extension to this model which assumes that the demand for some products will decline at the same rate. Using our assumption of Poisson demand, we obtained LTB quantities for target fill rate using a base stock policy.

Our forecasting approach was time-series based, using only previous demand observations to estimate model parameters. Times-series methodology to determine long-term forecasts for spare parts is under-researched and recent developments centre on short-term, opposed to long-term, decision-making. As such, our work makes an important contribution to the literature. In addition, the problem we solve has practical relevance. Our methodology aims to support planners who must make high-stake LTB decisions which are common in aftermarket industries. The forecasts determined using our approach could also be used for other long-term operational decisions including those regarding capacity and financial planning. Accurate long-term demand forecasts have economic and environmental advantages in a wide range of contexts.

We tested our methodology in both a simulated and real-world setting. We evaluated the performance of the forecasts on the total demand across the out-of-sample period and assessed the effectiveness of our approaches for LTBs. We measured the achieved fill rate with a target of 95% and calculated the excess inventory percentage to account for the difference in demand volumes across parts. As expected, we found that the forecasts for faster-moving products were more accurate than forecasts for slower-moving products. The results showed that our model extension provided increased forecast accuracy and returned a fill rate closer to the target while incurring less excess inventory.

We proposed that demand in each period be represented by a Poisson random variable. This is a common representation of the demand for aftermarket products in the literature (Teunter and Haneveld, 1998; Syntetos et al., 2012; Boylan and Syntetos, 2021). However, there are alternative distributional assumptions which could be made to model demand. Compound Poisson distributions offer more flexibility, as the variance can take a different value to the mean of demand. This is more suitable to model a series which has periods with large quantities and periods of zero demands. Another appropriate setting may be when demand quantities are large but the variance of these observations is small (Boylan and Syntetos, 2021). Further research might explore alternative distributional assumptions in modelling demand in the EOL phase. Furthermore, our approach assumes that demand observations are independent. Even though our model and its extension perform well both on simulated and real data a more sophisticated approach might incorporate auto-correlation of demand observations.

We used the demand history of all the products in the dataset when implementing the multiple series approach. In the example at the automotive manufacturer this works well. However, some inventories may have products with differing patterns of decline. Implementing our methodology, in this case, would necessitate techniques to group products with similar decline rates. To this aim, time-series-based clustering approaches that define proximity between two series based on the overall pattern of demand could be used. Alternatively, product information could be utilised. It would be interesting to determine which product attributes are most useful in the effective grouping of products with similar demand patterns.

A limitation of this work is that it does not consider issues in the integration of forecasting methodology and inventory decision-making. The literature on inventory management in almost all cases assumes that the demand distribution and its parameters are known to the decision-maker. In practice, decision-makers rely on estimated parameters and thus bias in the parameter estimates is propagated through to the calculation of the order quantities (Prak et al., 2017; Prak and Teunter, 2019). There is a shortage of work that considers the impact of forecast uncertainty on inventory control for demand which is Poisson and non-stationary. This is an important issue for future research.

5.2 Life Cycle Demand Forecasting

In Chapter 3, we focused on forecasts required throughout the product life cycle. It is especially challenging to determine a forecast over a product's life when demand has only been observed in the early phases of the life cycle. We introduce an approach that models demand over the life cycle as a Poisson process with an intensity function that follows the Bass curve for product adoption (Bass, 1969). We acknowledge widely discussed issues with parameter estimation for life cycle models when only a small amount of demand history has been observed (Srinivasan and Mason, 1986; Van den Bulte and Lilien, 1997; Putsis Jr and Srinfvasan, 2000; Meade and Islam, 2006) and thus extend our formulation by pooling the incomplete demand history from multiple series to estimate common parameters.

Our findings make an important contribution to the field of life cycle demand forecasting. Many of the applications of the Bass curve found in the literature assume that the demand for active products will be similar to past products. Thus, the full demand history of past products is used to forecast the demand for an active product over the life cycle. In our model extension, the incomplete demand histories of multiple active products were used to obtain joint parameters. This approach is advantageous for products for which manufacturers do not have full demand records, such as newly designed products and products with long life cycles.

Our model was validated using simulated data and real data from the automotive manufacturer. We calculated the percentage errors and absolute percentage errors for the total demand over the out-of-sample. Additionally, we determined the root mean squared scaled error (Makridakis et al., 2022) to account for scenarios where it is necessary to measure the forecast error in each period. In the simulation study, we showed that the mean absolute percentage error decreased as the number of series pooled for our model extension increased. Our experiment on 175 SKUs revealed that our model extension improves forecast accuracy irrespective of the length of the in-sample. This is especially notable for forecasting life cycle patterns when demand history is limited.

As in Chapter 2, our approach modelled demand using a common representation for aftermarket products and an assumption of independent demand observations. Alternative distributions suitable for the demand for spare parts include the Stuttering Poisson and the Negative Binomial distribution (Boylan and Syntetos, 2021). Furthermore, as demand sizes across spare parts can vary significantly, more than one distribution may need to be considered to model the demand of SKUs belonging to large inventories. Future research could explore alternative distributions and model formulations which assume dependence between demand observations.

In testing our model on real data we found commonality in the life cycle shape across the automotive spare parts. In other cases, assuming joint parameters may not be appropriate. Clustering could be used to group products with similar life cycle behaviour together prior to the implementation of our model. Additionally, we assume that the underlying intensity function driving demand follows the Bass curve. Polynomial, triangular and trapezoid life cycle shapes have also been discussed in the literature (Hu et al., 2019). It would be interesting to conduct an empirical investigation on the goodness-of-fit of various shapes for aftermarket products.

5.3 Outreach for Long-Term Impact

In Chapter 4 of this thesis, we explored the communication of forecasting and modelling topics to non-scientific audiences. The effective dissemination of forecasting and modelling techniques is important for many reasons. Firstly, the successful transition of knowledge ensures research developments lead to impact in the real world. Stakeholders must understand model formulations and methods to properly implement them within their systems. Further, OR methodology is commonly designed for real-world problems which impact the general population. Communication helps to build understanding between professionals and the wider public. Efforts to foster interest in statistics and OR also secure the recruitment of future professionals required to meet current workforce shortages and address diversity issues in STEM industries.

We showcase a diverse portfolio of outreach work based on forecasting and modelling topics designed based on core principles. We identified key contributors to the underrepresentation of girls pursuing STEM-related career pathways and designed content to oppose these factors. On a reflection of teacher feedback, we found that the content delivered to secondary school students was effective in developing skills and promoted interest in statistics and OR. This supported the view that forecasting and modelling techniques provide an exciting basis for outreach activity. We found that content which aimed at dismantling stereotypes related to the identity and competence of mathematical scientists was positively received.

This work mostly explores applications of forecasting and modelling methodology within television and film. A wider range of applications could enhance the appeal of this content. A benefit of our work is that we consider issues related to equity, diversity and inclusion. However, we primarily focus on issues related to the underrepresentation between women and girls. Future work that aims to overcome wider issues of inclusion in STEM are therefore suggested.

In a review of teacher feedback, we found that this work resulted in two students choosing to continue studying A-level mathematics when they had previously considered withdrawing from their course. This is a clear indication that our contribution has been effective in securing long-term impact. Nevertheless, longitudinal studies should be undertaken to better understand the impact of outreach work in achieving long-term goals. This work has laid the foundations for important future research in this area.

Bibliography

Altay, N., Rudisill, F., and Litteral, L. A. (2008). Adapting Wright's modification of Holt's method to forecasting intermittent demand. *International Journal of Production Economics*, 111(2):389–408.

Appel, D., Tillinghast, R. C., and Mansouri, M. (2021). Identifying positive catalysts in the stem career pipeline. In *2021 IEEE Integrated STEM Education Conference (ISEC)*, pages 132–139. IEEE.

Arenas-Vasco, A., Aguirre-González, E. J., Camacho, C., and Villegas, J. G. (2024). Among us as an operations research problem. *INFORMS Transactions on Education*, 24(3):282–284.

Bass, F. M. (1969). A new product growth for model consumer durables. *Management Science*, 15(5):215–227.

BBC News (2020). BTS Black Lives Matter: Fans match band's $1m donation. www.bbc.co.uk/news/world-asia-52960617. Accessed: 27-03-2025.

Behfard, S., Al Hanbali, A., Van der Heijden, M. C., and Zijm, W. H. (2018). Last time buy and repair decisions for fast moving parts. *International Journal of Production Economics*, 197:158–173.

Behfard, S., Van der Heijden, M. C., Al Hanbali, A., and Zijm, W. H. (2015). Last time buy and repair decisions for spare parts. *European Journal of Operational Research*, 244(2):498–510.

Bellman, R. and Kalaba, R. (1959). On adaptive control processes. *IRE Transactions on Automatic Control*, 4(2):1–9.

Bellos, A. (2011). *Alex's Adventures in Numberland*. A&C Black.

Bertrand, M. and Duflo, E. (2017). Field experiments on discrimination. *Handbook of Economic Field Experiments*, 1:309–393.

Box Office Mojo (2019). Knives Out. www.boxofficemojo.com/release/rl3204875777/. Accessed: 27-03-2025.

Boylan, J. E. and Syntetos, A. A. (2021). *Intermittent Demand Forecasting: Context, Methods and Applications*. John Wiley & Sons.

Bradshaw, N.-A. and Mann, T. (2021). Welcoming students to the mathematics community: obstacles to "belonging". *MSOR Connections*, 19(2):75–86.

Brams, S. J. (1993). Theory of Moves. *American Scientist*, 81(6):562–570.

Brams, S. J. (2011). *Game Theory and the Humanities: Bridging Two Worlds*. MIT press.

British Science Association (2020). The State of the Sector: Diversity and representation in STEM industries in the UK. www.britishscienceassociation.org/Handlers/Download.ashx?IDMF=d7899dce-22d5-4880-bbcf-669c0c35bda6. Accessed: 27-03-2025.

Bruce, L. (2022). The secret genius of modern life. *BBC Studios Science Unit*.

Budd, C. J. (2015). Promoting maths to the general public. In *The Oxford Handbook of Numerical Cognition*, pages 3–16. Oxford University Press.

Bull, A. (2024). Addressing sexual misconduct in higher education, part one: Prevention. www.timeshighereducation.com/campus/addressing-sexual-misconduct-higher-education-part-one-prevention. Accessed: 27-03-2025.

Campaign for Science and Engineering (CaSE) (2023). The Skills Opportunity Report. www.sciencecampaign.org.uk/app/uploads/2023/06/CaSE-The-Skills-Opportunity-Report-June-2023.pdf. Accessed: 27-03-2025.

Carroll, P. and Esposito Amideo, A. (2024). Gender equality: opportunities and challenges for the OR community. *Journal of the Operational Research Society*, 75(6):1011–1029.

Cassidy, R., Cattan, S., Crawford, C., and Dytham, S. (2018). How can we increase girls' uptake of maths and physics A-level? Technical report, The Institute for Fiscal Studies.

Chamberlain, N. (2020). The Black Heroes of Mathematics. www.youtube.com/watch?v=YuYVa6aDFRE. Accessed: 27-03-2025.

Chan, T. C., Fernandes, C., Loa, A., and Sandholtz, N. (2024). Case article—moneyball for murderball: Using analytics to construct lineups in wheelchair rugby. *INFORMS Transactions on Education*, 24(2):175–181.

Chandrasekaran, D. and Tellis, G. J. (2007). A critical review of marketing research on diffusion of new products. In *Review of Marketing Research*, pages 39–80. Emerald Group Publishing Limited.

Cheryan, S., Master, A., and Meltzoff, A. N. (2015). Cultural stereotypes as gatekeepers: Increasing girls' interest in computer science and engineering by diversifying stereotypes. *Frontiers in Psychology*, 6:123074.

Chin, C. (2022). BTS make history with the most Number One hits on the Billboard Hot 100 this decade. www.nme.com/news/music/bts-most-number-one-hits-billboard-hot-100-chart-decade-3341793. Accessed: 27-03-2025.

Chlond, M. J. (2016). TSP at the movies: Yondu's dart problem. *INFORMS Transactions on Education*, 16(3):110–112.

Chou, Y.-C., Hsu, Y. S., and Lu, S.-Y. (2016). A demand forecast method for the final ordering problem of service parts. *International Journal of Industrial Engineering*, 23(2).

Clark, S. L., Dyar, C., Inman, E. M., Maung, N., and London, B. (2021). Women's career confidence in a fixed, sexist stem environment. *International Journal of STEM Education*, 8:1–10.

Cochran, J. J. (2004). Introduction to the special issue: SpORts in the OR classroom. *INFORMS Transactions on Education*, 5(1):1–1.

Cohen, M. C., Zhang, R., and Jiao, K. (2022). Data aggregation and demand prediction. *Operations Research*, 70(5):2597–2618.

Colajanni, G., Gobbi, A., Picchi, M., Raffaele, A., and Taranto, E. (2023). An Operations Research–based teaching unit for grade 10: The ROAR experience, part i. *INFORMS Transactions on Education*, 23(2):104–120.

Crawford, A., Hays, C. L., Schlichte, S. L., Greer, S. E., Mallard, H. J., Singh, R. M., Clarke, M. A., and Schiller, A. M. (2021). Retrospective analysis of a stem outreach event reveals positive influences on student attitudes toward stem careers but not scientific methodology. *Advances in Physiology Education*, 45(3):427–436.

Cronin, A., Ní Shúilleabháin, A., Lewanowski-Breen, E., and Kennedy, C. (2017). Maths sparks: investigating the impact of outreach on pupils' attitudes towards mathematics. *MSOR Connect*, 15:4–13.

Croston, J. D. (1972). Forecasting and stock control for intermittent demands. *Journal of the Operational Research Society*, 23(3):289–303.

David, B. (2014). Introduction: Making the move to peer learning. In *Peer Learning in Higher Education*, pages 1–17. Routledge.

Dekker, R., Pinçe, Ç., Zuidwijk, R., and Jalil, M. N. (2013). On the use of installed base information for spare parts logistics: A review of ideas and industry practice. *International Journal of Production Economics*, 143(2):536–545.

Dellarocas, C., Zhang, X. M., and Awad, N. F. (2007). Exploring the value of online product reviews in forecasting sales: The case of motion pictures. *Journal of Interactive Marketing*, 21(4):23–45.

Department for Business, Innovation & Skills (2012). The 2011 skills for life survey: A survey of literacy, numeracy and ICT levels in England. www.assets.publishing.service.gov.uk/media/5a74946f40f0b61938c7e98f/12-p168-2011-skills-for-life-survey.pdf. Accessed: 27-03-2025.

DePuy, G. W. and Taylor, G. D. (2007). Using board puzzles to teach Operations Research. *INFORMS Transactions on Education*, 7(2):160–171.

Dombi, J., Jónás, T., and Tóth, Z. E. (2018). Modeling and long-term forecasting demand in spare parts logistics businesses. *International Journal of Production Economics*, 201:1–17.

Durugbo, C. M. (2020). After-sales services and aftermarket support: a systematic review, theory and future research directions. *International Journal of Production Research*, 58(6):1857–1892.

Elandt-Johnson, R. C. and Johnson, N. L. (1980). *Survival Models and Data Analysis*. John Wiley & Sons, 1st edition.

Fortuin, L. (1980). The all-time requirement of spare parts for service after sales—theoretical analysis and practical results. *International Journal of Operations & Production Management*, 1(1):59–70.

Fry, H. (2018). *Hello World: How to be Human in the Age of the Machine*. Random House.

Ganjeizadeh, F., Lei, H., Goraya, P., and Olivar, E. (2017). Applying looks-like analysis and bass diffusion model techniques to forecast a neurostimulator device with no historical data. *Procedia Manufacturing*, 11:1916–1924.

Glimpse (2022). Glimpse – Google Trends supercharged. www.meetglimpse.com/extension/. Accessed: 27-03-2025.

González-Pérez, S., Mateos de Cabo, R., and Sáinz, M. (2020). Girls in STEM: Is it a female role-model thing? *Frontiers in Psychology*, 11:564148.

Good, C., Rattan, A., and Dweck, C. S. (2012). Why do women opt out? Sense of belonging and women's representation in mathematics. *Journal of Personality and Social Psychology*, 102(4):700.

Goossens, D. and Beliën, J. (2023). Teaching integer programming by scheduling the belgian soccer league. *INFORMS Transactions on Education*, 23(3):164–172.

Greenfield, D. (2007). What's the deal with the white middle-aged guy teaching hip-hop? lessons in popular culture, positionality and pedagogy. *Pedagogy, Culture & Society*, 15(2):229–243.

Guidolin, M. and Manfredi, P. (2023). Innovation diffusion processes: Concepts, models, and predictions. *Annual Review of Statistics and Its Application*, 10(1):451–473.

Hartmann, S. (2018). Puzzle—solving smartphone puzzle apps by mathematical programming. *INFORMS Transactions on Education*, 18(2):127–141.

Hasni, M., Aguir, M., Babai, M., and Jemai, Z. (2019). Spare parts demand forecasting: a review on bootstrapping methods. *International Journal of Production Research*, 57(15-16):4791–4804.

Hauser, J., Tellis, G. J., and Griffin, A. (2006). Research on innovation: A review and agenda for marketing science. *Marketing Science*, 25(6):687–717.

Holt, C. (1957). Forecasting seasonals and trends by exponentially weighted averages (onr memorandum no. 52). *Carnegie Institute of Technology, Pittsburgh USA*, 10.

Hong, J. S., Koo, H.-Y., Lee, C.-S., and Ahn, J. (2008). Forecasting service parts demand for a discontinued product. *IIE Transactions*, 40(7):640–649.

Hu, K., Acimovic, J., Erize, F., Thomas, D. J., and Van Mieghem, J. A. (2019). Forecasting new product life cycle curves: Practical approach and empirical analysis. *Manufacturing & Service Operations Management*, 21(1):66–85.

Imafidon, A.-M. (2022). *She's in CTRL: How Women Can Take Back Tech*. Bantam Press.

Inderfurth, K. and Kleber, R. (2013). An advanced heuristic for multiple-option spare parts procurement after end-of-production. *Production and Operations Management*, 22(1):54–70.

Ismail, Z. and Abu, N. (2013). New car demand modeling and forecasting using bass diffusion model. *American Journal of Applied Sciences*, 10(6):536–541.

Johnson, A. and Mulligan, C. (2016). What makes a successful outreach kit? *MSOR Connections*, 14(2):28–33.

Johnston, F., Boylan, J. E., and Shale, E. A. (2003). An examination of the size of orders from customers, their characterisation and the implications for inventory control of slow moving items. *Journal of the Operational Research Society*, 54:833–837.

Kaiser, G., Bracke, M., Göttlich, S., and Kaland, C. (2013). Authentic complex modelling problems in mathematics education. *Educational Interfaces between Mathematics and Industry: Report on an ICMI-ICIAM-Study*, pages 287–297.

Keeling, M. J. and Rohani, P. (2008). Introduction to simple epidemic models. In *Modeling Infectious Diseases in Humans and Animals*, pages 15–53. Princeton University Press.

Kim, R. (2022a). *Hallyu!: The Korean Wave.* V&A Publications, 1st edition.

Kim, R. (2022b). Hallyu! The Korean Wave. Exhibition, 24 September 2022 - 25 June 2023, Victoria and Albert Museum, South Kensington, London.

Kim, T., Hong, J., and Kang, P. (2015). Box office forecasting using machine learning algorithms based on SNS data. *International Journal of Forecasting*, 31(2):364–390.

Kim, T. Y., Dekker, R., and Heij, C. (2017). Spare part demand forecasting for consumer goods using installed base information. *Computers & Industrial Engineering*, 103:201–215.

Kolassa, S. (2016). Evaluating predictive count data distributions in retail sales forecasting. *International Journal of Forecasting*, 32(3):788–803.

Koopman, C. W. (2011). Optimizing the last time buy decision at the IBM service part operation organization. Master's thesis, University of Twente.

Kruskal, J. B. (1983). An overview of sequence comparison: Time warps, string edits, and macromolecules. *SIAM review*, 25(2):201–237.

Lancaster University (2023a). The Florence Nightingale Days. wwww.lancaster.ac.uk/maths/engagement/working-with-schools/florence-nightingale-day/january-2023/. Accessed: 27-03-2025.

Lancaster University (2023b). Year 12 STEM Taster Days. www.lancaster.ac.uk/sci-tech/about-us/schools-outreach/year-12-stem-taster-day/. Accessed: 27-03-2025.

Lathisms (2016). Latinxs and Hispanics in the Mathematical Sciences. www.lathisms.org/. Accessed: 27-03-2025.

Lei, D., Qi, Y., Liu, S., Geng, D., Zhang, J., Hu, H., and Shen, Z.-J. M. (2024). Pooling and boosting for demand prediction in retail: A transfer learning approach. *Manufacturing & Service Operations Management*.

Li, X., Yin, Y., Manrique, D. V., and Bäck, T. (2021). Lifecycle forecast for consumer technology products with limited sales data. *International Journal of Production Economics*, 239:108206.

Li, Y. and Sui, M. (2011). Literature analysis of innovation diffusion. *Technology and Investment*, 2(3):155–162.

Lilien, G. L., Rangaswamy, A., and De Bruyn, A. (2017). *Principles of Marketing Engineering and Analytics 3rd Edition*. DecisionPro.

Linke, I. and Hunsicker, E. (2018). Faces of Women in Mathematics. www.vimeo.com/259039018. Accessed: 27-03-2025.

London Mathematics Society (2023). Gender in mathematics – quantitative analysis. www.lms.ac.uk/sites/default/files/inline-files/PN2113_LMS_Benchmarking_March2023_v0-4_3.pdf. Accessed: 27-03-2025.

Makridakis, S., Spiliotis, E., and Assimakopoulos, V. (2022). M5 accuracy competition: Results, findings, and conclusions. *International Journal of Forecasting*, 38(4):1346–1364.

Master, A., Meltzoff, A. N., and Cheryan, S. (2021). Gender stereotypes about interests start early and cause gender disparities in computer science and engineering. *Proceedings of the National Academy of Sciences*, 118(48):e2100030118.

Mathematics Education Innovation (MEI) (2018). Advanced mathematics support programme. www.amsp.org.uk. Accessed: 27-03-2025.

Meade, N. and Islam, T. (2006). Modelling and forecasting the diffusion of innovation–a 25-year review. *International Journal of forecasting*, 22(3):519–545.

Mehringer, J. and Menden, C. (2019). Forecasting spare parts demand with clustering approaches and functional approximations. Working paper, Fraunhofer Center for Applied Research on Supply Chain Services & University of Erlangen-Nuremberg.

Mercer-Mapstone, L. and Kuchel, L. (2017). Core skills for effective science communication: A teaching resource for undergraduate science education. *International Journal of Science Education, Part B*, 7(2):181–201.

Moore, J. R. (1971). Forecasting and scheduling for past-model replacement parts. *Management Science*, 18(4-part-i):B–200.

Osborne, M. J. (2004). An Introduction to Game Theory. *Oxford University Press*, 2:672–713.

Pagel, C. (2023). More cancer screenings sounds like good news. But can the NHS cope with them? www.theguardian.com/commentisfree/2023/aug/28/cancer-screenings-nhs-detection-treatment-health-service. Accessed: 27-03-2025.

Pawitan, Y. (2001). *In All Likelihood: Statistical Modelling and Inference Using Likelihood*. Oxford University Press.

Peres, R., Muller, E., and Mahajan, V. (2010). Innovation diffusion and new product growth models: A critical review and research directions. *International Journal of Research in Marketing*, 27(2):91–106.

Prak, D. and Teunter, R. (2019). A general method for addressing forecasting uncertainty in inventory models. *International Journal of Forecasting*, 35(1):224–238.

Prak, D., Teunter, R., and Syntetos, A. (2017). On the calculation of safety stocks when demand is forecasted. *European Journal of Operational Research*, 256(2):454–461.

Prestwich, S. D., Tarim, S. A., and Rossi, R. (2021). Intermittency and obsolescence: A croston method with linear decay. *International Journal of Forecasting*, 37(2):708–715.

Prestwich, S. D., Tarim, S. A., Rossi, R., and Hnich, B. (2014). Forecasting intermittent demand by hyperbolic-exponential smoothing. *International Journal of Forecasting*, 30(4):928–933.

Putsis Jr, W. P. and Srinfvasan, V. (2000). Estimation technqiues for macro diffusion models. In *New-Product Diffusion Models*, chapter 11, pages 263–294. Springer Science & Business Media.

Raffaele, A. and Gobbi, A. (2021). Teaching Operations Research before university: A focus on grades 9–12. In *Operations Research Forum*, volume 2, page 13. Springer.

Rakedzon, T. (2019). The challenges of writing science: Tools for teaching and assessing written science communication. In *Theory and Best Practices in Science Communication Training*, pages 90–103. Routledge.

Reinking, A. and Martin, B. (2018). The gender gap in STEM fields: Theories, movements, and ideas to engage girls in STEM. *Journal of New Approaches in Educational Research*, 7(2):148–153.

Reis, K. and Wadkin, L. (2023). Women & gender-minorities in mathematics: PhD student challenges. www.staff.ncl.ac.uk/laurawadkin/files/2023/03/LWadkin_MathematicsPhDStudentExperience.pdf. Accessed: 27-03-2025.

Ritchie, E. and Wilcox, P. (1977). Renewal theory forecasting for stock control. *European Journal of Operational Research*, 1(2):90–93.

Roll, M. (2021). Korean wave (hallyu) - the rise of Korea's cultural economy & pop culture. www.martinroll.com/resources/articles/asia/korean-wave-hallyu-the-rise-of-koreas-cultural-economy-pop-culture. Accessed: 27-03-2025.

Royal Instituition (2024). Ri Masterclasses. www.rigb.org/learning/ri-masterclasses. Accessed: 27-03-2025.

Royal Society (2021). Science education for a research and innovation economy. www.royalsociety.org/news-resources/publications/2022/science-education-for-a-research-and-innovation-economy/. Accessed: 27-03-2025.

Rutherford, A. (2018). *A brief history of everyone who ever lived: The human story retold through our genes*. The Experiment.

Sáinz, M., Fàbregues, S., Romano, M. J., and López, B.-S. (2022). Interventions to increase young people's interest in STEM. a scoping review. *Frontiers in Psychology*, 13:954996.

Sanguri, K., Patra, S., Nikolopoulos, K., and Punia, S. (2024). Intermittent demand, inventory obsolescence, and temporal aggregation forecasts. *International Journal of Production Research*, 62(5):1663–1685.

Sarlo, R., Fernandes, C., and Borenstein, D. (2023). Lumpy and intermittent retail demand forecasts with score-driven models. *European Journal of Operational Research*, 307(3):1146–1160.

Saunders, N. (2024). Maths degrees are becoming less accessible – and this is a problem for business, government and innovation. www.theconversation.com/maths-degrees-are-becoming-less-accessible-and-this-is-a-problem-for-business-government-and-innovation-227488. Accessed: 27-03-2025.

Science and Technology Committee (2023). Inquiry into diversity and inclusion in STEM. www.publications.parliament.uk/pa/cm5803/cmselect/cmsctech/95/report.html. Accessed: 27-03-2025.

Semenoglou, A.-A., Spiliotis, E., Makridakis, S., and Assimakopoulos, V. (2021). Investigating the accuracy of cross-learning time series forecasting methods. *International Journal of Forecasting*, 37(3):1072–1084.

Senin, P. (2008). Dynamic time warping algorithm review. www.seninp.github.io/assets/pubs/senin_dtw_litreview_2008.pdf. Accessed: 27-03-2025.

Shim, D. (2006). Hybridity and the rise of Korean popular culture in Asia. *Media, Culture & Society*, 28(1):25–44.

Shin, J. E. L., Levy, S. R., and London, B. (2016). Effects of role model exposure on STEM and non-STEM student engagement. *Journal of Applied Social Psychology*, 46(7):410–427.

Sillanpää, V. and Liesiö, J. (2018). Forecasting replenishment orders in retail: value of modelling low and intermittent consumer demand with distributions. *International Journal of Production Research*, 56(12):4168–4185.

Simon, G., Lendasse, A., Cottrell, M., Fort, J.-C., and Verleysen, M. (2005). Time series forecasting: Obtaining long term trends with self-organizing maps. *Pattern Recognition Letters*, 26(12):1795–1808.

Snyder, R. D., Ord, J. K., and Beaumont, A. (2012). Forecasting the intermittent demand for slow-moving inventories: A modelling approach. *International Journal of Forecasting*, 28(2):485–496.

Spangler, T. (2021). Squid Game is decisively Netflix no. 1 show of all time with 1.65 billion hours streamed in first four weeks, company says. www.variety.com/2021/digital/news/squid-game-all-time-most-popular-show-netflix-1235113196/. Accessed: 27-03-2025.

Spooner, K., Nomani, J., and Cook, S. (2024). Improving high school students' perceptions of mathematics through a mathematical modelling course. *Teaching Mathematics and its Applications: An International Journal of the IMA*, 43(1):38–50.

Srinivasan, V. and Mason, C. H. (1986). Nonlinear least squares estimation of new product diffusion models. *Marketing Science*, 5(2):169–178.

Starr, C. R. (2018). "I'm not a science nerd!" STEM stereotypes, identity, and motivation among undergraduate women. *Psychology of Women Quarterly*, 42(4):489–503.

Steckles, K. (2025). Barcodes: How they could be your latest mathematical party trick. www.newscientist.com/article/mg26535291-000-barcodes-how-they-could-be-your-latest-mathematical-party-trick/. Accessed: 27-03-2025.

Steckles, K., Rowlett, P., and Ugonna, A. (2020). Pre-university informal engagement with mathematical activities and the decision to study mathematics at university. *MSOR Connections*, 18(3):10–22.

STEM Women (2021). LGBTQ+ in STEM. www.stemwomen.com/lgbtq-plus-in-stem. Accessed: 27-03-2025.

Stuart, A. and Ord, J. K. (1998). *Kendall's Advanced Theory of Statistics*. Arnold, 6th edition.

Sullivan Jr, W. J., Strathdee, S. A., and Hotez, P. J. (2021). Making it personal: science communication for the masses. *Trends in Parasitology*, 37(8):684–686.

Syntetos, A., Lengu, D., and Babai, M. Z. (2013). A note on the demand distributions of spare parts. *International Journal of Production Research*, 51(21):6356–6358.

Syntetos, A. A., Babai, M. Z., and Altay, N. (2012). On the demand distributions of spare parts. *International Journal of Production Research*, 50(8):2101–2117.

Syntetos, A. A., Babai, Z., Boylan, J. E., Kolassa, S., and Nikolopoulos, K. (2016). Supply chain forecasting: Theory, practice, their gap and the future. *European Journal of Operational Research*, 252(1):1–26.

Syntetos, A. A. and Boylan, J. E. (2005). The accuracy of intermittent demand estimates. *International Journal of Forecasting*, 21(2):303–314.

Syntetos, A. A., Boylan, J. E., and Croston, J. (2005). On the categorization of demand patterns. *Journal of the Operational Research Society*, 56(5):495–503.

Taranto, E., Colajanni, G., Gobbi, A., Picchi, M., and Raffaele, A. (2024). Fostering students' modelling and problem-solving skills through Operations Research, digital technologies and collaborative learning. *International Journal of Mathematical Education in Science and Technology*, 55(8):1957–1998.

Teunter, R. H. and Fortuin, L. (1999). End-of-life service. *International Journal of Production Economics*, 59(1-3):487–497.

Teunter, R. H. and Haneveld, W. K. K. (1998). The 'final order' problem. *European Journal of Operational Research*, 107(1):35–44.

Teunter, R. H., Syntetos, A. A., and Babai, M. Z. (2011). Intermittent demand: Linking forecasting to inventory obsolescence. *European Journal of Operational Research*, 214(3):606–615.

The Movie Database (2021). www.themoviedb.org. Accessed: 27-03-2025.

Thompson, G., Hawkins, O., Dar, A., and Taylor, M. (2012). *Olympic Britain: social and economic changes since the 1908 and 1948 London games*. House of Commons Library.

Tillinghast, R. C., Appel, D. C., Winsor, C., and Mansouri, M. (2020). Stem outreach: A literature review and definition. In *2020 IEEE Integrated STEM Education Conference (ISEC)*, pages 1–20. IEEE.

University of Greenwich (2023). Greenwich Maths Time. www.gre.ac.uk/articles/engsci/greenwich-maths-time-the-festival-of-mathematics-and-its-applications. Accessed: 27-03-2025.

Van Camp, A. R., Gilbert, P. N., and O'Brien, L. T. (2019). Testing the effects of a role model intervention on women's stem outcomes. *Social Psychology of Education*, 22:649–671.

Van den Bulte, C. and Lilien, G. L. (1997). Bias and systematic change in the parameter estimates of macro-level diffusion models. *Marketing Science*, 16(4):338–353.

Van der Auweraer, S. and Boute, R. (2019). Forecasting spare part demand using service maintenance information. *International Journal of Production Economics*, 213:138–149.

Van der Auweraer, S., Boute, R. N., and Syntetos, A. A. (2019). Forecasting spare part demand with installed base information: A review. *International Journal of Forecasting*, 35(1):181–196.

Van der Heijden, M. and Iskandar, B. P. (2013). Last time buy decisions for products sold under warranty. *European Journal of Operational Research*, 224(2):302–312.

Venema, V., Lee, J. Y., and Haque, S. (2023). K-drama: The women pushing boundaries on TV. www.bbc.co.uk/news/world-asia-67583282. Accessed: 27-03-2025.

Vennix, J., Den Brok, P., and Taconis, R. (2018). Do outreach activities in secondary stem education motivate students and improve their attitudes towards stem? *International Journal of Science Education*, 40(11):1263–1283.

Weale, S. (2023). Shortage of teachers will be a big maths problem for Rishi Sunak. www.theguardian.com/education/2023/apr/17/shortage-of-teachers-will-be-a-big-maths-problem-for-rishi-sunak. Accessed: 27-03-2025.

Weale, S. (2024). Disadvantaged parents in England and Wales rule out STEM jobs for children, charity says. www.theguardian.com/education/2024/jan/23/disadvantaged-parents-england-wales-rule-out-stem-jobs-for-children-charity-says. Accessed: 27-03-2025.

WISE (2019). 2019 workforce statistics – one million women in STEM in the UK. www.wisecampaign.org.uk/2019-workforce-statistics-one-million-women-in-stem-in-the-uk/. Accessed: 27-03-2025.

Womanthology (2024). Connecting women and opportunity. www.womanthology.co.uk/. Accessed: 27-03-2025.

Wright, D. J. (1986). Forecasting data published at irregular time intervals using an extension of Holt's method. *Management Science*, 32(4):499–510.

Yates, K. (2019). *The Maths of Life and Death: 7 Mathematical Principles That Shape Our Lives*. Hachette UK.

Yates, K. (2023). Has the time come to stop changing the clocks? This is what the science tells us. www.theguardian.com/commentisfree/2023/oct/28/eu-daylight-saving-british-summer-time-greenwich-mean-. Accessed: 27-03-2025.

Zhuhadar, L. P. (2020). SIR modeling. https://people.wku.edu/lily.popova.zhuhadar/. Accessed: 27-03-2025.

Zipkin, P. (2000). *Foundations of Inventory Management*. McGraw-Hill Companies, Incorporated.