




INFORMS Transactions on Education

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Case—Canyon Bicycles: Judgmental Demand Forecasting in Direct Sales

Christoph Diermann, Arnd Huchzermeier

To cite this article:

Christoph Diermann, Arnd Huchzermeier (2017) Case—Canyon Bicycles: Judgmental Demand Forecasting in Direct Sales. INFORMS Transactions on Education 17(2):63-74. <https://doi.org/10.1287/ited.2016.0165cs>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2017, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Case

Canyon Bicycles: Judgmental Demand Forecasting in Direct Sales

Christoph Diermann,^a Arnd Huchzermeier^a

^a WHU—Otto Beisheim School of Management, 56179 Vallendar, Germany

Contact: christoph.diermann@whu.edu (CD); arnd.huchzermeier@whu.edu (AH)

Received: March 4, 2016

Accepted: July 8, 2016

Published Online: January 13, 2017

Supplemental Material: Supplemental material is available at <https://doi.org/10.1287/ited.2016.0165cs>.

<https://doi.org/10.1287/ited.2016.0165cs>

Copyright: © 2017 INFORMS

Keywords: judgmental demand forecasting • demand uncertainty • debiasing demand forecasts • forecast accuracy

1. Introduction

One beautiful morning in March of 2013, Mr. Roman Arnold, CEO of the family owned company, Canyon Bicycles, was in deep thought. He went up the stairs to his home office to prepare for one of the most critical tasks facing Canyon Bicycles. During the next two days, he and a group of company experts would—as they did every year—forecast and commit to fulfilling the next season's entire product demand, thereby determining a significant share of the season's success. Forecasting was a pressing issue at Canyon Bicycles: product portfolios had continued to widen even as product life cycles (along with customer patience regarding delivery) were becoming shorter. As always, the company's order commitment to its suppliers had to be made six months *before* the season start in September, that is, without having any early indication of demand. Although his management team consisted of strongly committed employees, Mr. Arnold was aware that the team's consensus demand forecasts had in past years been consistently too low. Canyon Bicycles had grown rapidly in recent years and along with it the firm's financing risk, since all major components had to be prefinanced by a local savings bank. To address the team's demand forecasting problem, Mr. Arnold had sought advice from operations management experts at the nearby business school. He expressed an urgent need to improve the management team's forecasting process in order to increase the firm's profitability and mitigate financing risks. Over the next two days, the newly developed forecasting approach would be applied for the first time.

The history of Canyon Bicycles, an innovative premium bike manufacturer and online retailer located in Koblenz, Germany, goes back to 1983. Roman Arnold, his father, and his brother, Lothar—who joined

the company as COO in 2015—were all passionate road bike cyclists who participated in many Italian road races. On the way from Germany to Italy, they always passed by various premium bike accessory manufacturers from which cyclists had to order when equipping themselves for the races. After observing other cyclists' interest in high-quality equipment, the Arnolds began to sell bike parts to them directly from a small trailer at the racetrack. The trio soon became so well known for their service that orders, to be delivered at the next race, were placed in advance at the current racetrack. Two years later, the company *Radsport Arnold* was officially founded. Despite being only 18 years old at the time, Roman Arnold assumed management of the start-up company, continuing the business as a small shop offering bike accessories and service in Koblenz.

In the mid-1990s, the company started to develop and manufacture its own high-end mountain bikes (MTB) and road bikes (RB); see Appendix A.1. In 2001, the company was renamed *Canyon Bicycles* and became a strong global brand for premium bikes; its product always ranks highly in tests and has won many design awards and readers' choice awards (see Appendix B). Each year in September, the Eurobike Fair in Friedrichshafen at Lake Constance, Germany, showcases the latest bike technology, and Canyon Bicycles is acclaimed as one of the industry's dominant design and innovation leaders. In particular, the company specializes in lightweight carbon fiber (CF) and aluminum (AL) frames. The venture's founding principle—of being close to the customer—has carried on and is now manifested in Canyon Bicycles' online direct sales strategy; see Appendix C.

Canyon Bicycles has exhibited exceptional growth over the years. Most customers were wealthy indi-

viduals who owned several premium bikes or young bike enthusiasts traveling and biking the world. Demand was highly seasonal, and the number of units sold increased, on average, by 23% annually between 2010 and 2014. The average sales price of a bike was about €2,900 and ranged from €600 to €7,800. In 2014, Canyon Bicycles invested €19.5 million in building a nearby production facility to serve the increasing demand; see Appendix D. The company employed more than 600 people (including temporary workers) in 2014 and shipped more than 70,000 bikes around the world that year. Total weekly demand across all bikes followed a strongly similar order pattern across different selling seasons (see Figure E.1).

Canyon Bicycles focuses on two major bike types: mountain bikes and road bikes. This distinction is based not only on their different design and uses but also on their respective demand patterns. Whereas mountain bikes are typically in demand throughout the year, road bikes are usually ordered during the good-weather season (see Figure E.2).

At a disaggregated level (i.e., item/bike) level, as illustrated in Appendix F, cumulative weekly demand of bikes within the same sales season did not follow a similar sales pattern (see Figure G.1, left). Also, across two sales seasons the weekly demand varied substantially, even for the same bike (see Figure G.1, right). These observations of past trends made it clear that early demand information could not in this case be used for (say) Bayesian demand forecast updating, which is often employed to estimate the future demand of fashion items.

The recent rapid growth of Canyon Bicycles has come at the cost of increased demand uncertainty accompanied by reduced forecast accuracy. Mr. Arnold had been relying on a team-based consensus forecast to determine next season's product demand for each bike. To accommodate the high growth, the forecasting team simply inflated the previous season's order quantities. Yet because the team forecasters could not anticipate whether a particular bike would be a fast- or a slow-moving item, they inflated orders for all bikes by the same factor.¹ To curtail excess inventory and reduce demands on the firm's tied-up capital, forecasters inflated the previous values by *less* than the expected growth rate because orders, once placed, could not be canceled. The resulting consistent underestimation of demand led to frequent stockouts and disappointed customers. However, this "conservative" ordering behavior was strongly reinforced by the market exits of competitors—for example, Fischer in 2009, E-Motion in 2011, MZ in 2012, Teikotec in 2012, PG in 2013, Mifa in 2014, and the famous Kettler in 2015.

Each year, Canyon Bicycles engaged in intense negotiations with its key financial institute to prefinance production by way of external borrowing. The

increased product demand uncertainty translated into increased financial risk, raising the costs of external capital. Both Mr. Arnold and the local savings bank knew that simply inflating all order quantities was not a sustainable growth strategy and could not be financed even in the medium term. Mr. Arnold recognized the importance of accurate demand forecasts not only for reducing supply demand mismatch costs but also for minimizing the restrictions and lowering the risk premium on funds borrowed by the company.

2. Product Development and Supply Chain Design

2.1. Development

Canyon Bicycles offered a wide variety of different aluminum and carbon fiber frames for their MTBs and RBs. The development of all frames and patent-protected manufacturing technologies was entirely organized in-house in the German headquarters in Koblenz. Any new product development required many steps and processes of different divisions to be run in parallel. Most design development, prototyping, and tests for compatibility, endurance, and function consisted of processes that were iterated until the frame's prespecified requirements were met, enabling clearance for production. Development time was driven mainly by product complexity. The development of an entirely new, wind-tunnel-optimized carbon road bike frame could take as long as 36 months, whereas improvements to an existing frame might require less than a year. The average development time for a new frame was about 18 months.

External factors, such as changes in the regulations of bike specifications for international races, can have a major effect on the development process and time. For instance, numerous decisions by the Union Cycliste Internationale (UCI) required that planned developments be sped up, altered, or canceled.

2.2. Sourcing and Lead Times

Although Canyon Bicycles designed, developed, and engineered their own frames, all production was outsourced to suppliers in Taiwan and China. Hence lead times were significant: on average, between 120 and 150 days passed between the firm's placing an order and receiving the goods. At least a third of that time was due solely to the (overseas) shipment process.

Downward volume adjustments to initial order quantities were seldom possible and never within four months of the season's start. During the sales season, order adjustments would not be effective because of the long delivery lead times. In urgent cases, the shipment time could be reduced to 10 days by using airfreight. However, that shipping option did little to improve the bikes' overall availability since orders were produced and shipped in fixed, predetermined batch sizes and

time intervals. Airfreight thus merely shifted the supply demand mismatch to a later time in the sales season. Upward adjustments in the order quantity were fraught with uncertainty and also depended on the supplier's idle processing capacity. In sum, the ordering decision was mostly a "one-shot" game for Canyon Bicycles.

Besides the frames, Canyon Bicycles also outsourced such bike components as wheels, brakes, gearing, suspension forks, chains, and cable hoists. Most components were sourced from suppliers in Europe, the United States, and Japan. The supply lead times for these parts varied between 60 and 90 days.

2.3. Assembly and Delivery

Once a customer placed an order for a bike online (or via the call center), the so-called picking list of the required parts would be forwarded on a rack to the shop floor. About 80% of all orders were placed online and the rest were placed via the call center. A local showroom contiguous to the factory only made a marginal contribution to sales. On average, three hours were needed to pick all parts and to prepare them for assembly; workers typically required another 2.5 hours for a bike's actual assembly. All bikes were tested after final assembly. Once all quality checks were passed, each bike was partly deconstructed and stored in the *Bike Guard*—the company's patented box for the shipment of their product.

Shipment from Canyon Bicycles to the customer required only one day within Germany and five days worldwide. Once the *Bike Guard* had arrived at the customer's home, just a few major parts had to be reassembled. Each bike came with a detailed manual that included directions for adjusting the bike to fit the customer's physique; the entire setup could be completed in only 10 minutes. The *Bike Guard* was designed not only for the bike's initial shipment to the customer but also for its later transportation. Examples include sending the bike back to Canyon Bicycles for service or transporting it aboard an airplane to such biking destinations as the Spanish island of Mallorca, which is visited by more than 100,000 bicycle enthusiasts each year.

3. The New Demand Forecasting Process

3.1. Measuring Demand

To gauge the accuracy of past demand forecasts, one need only compare them to the realized (actual) demand. At first it seemed a trivial matter to determine actual demand during the past sales season because all bikes were sold via channels (online and call center) that tracked customer preferences. Yet even this information can reveal the actual demand only for those bikes that, throughout the entire sales season, had never been out of stock. Thus the true demand

for bikes was unknown whenever demand exceeded supply. The online reporting of stockouts and expected lead times may well have influenced ordering decisions. Did customers switch to a different model when they found out that their bike of interest was out of stock? Did they wait for their preferred model to be back in stock? Did such customers buy a comparable bike from a competitor? Or did they end up not buying a bike at all?

Canyon resolved this ambiguity by defining actual season demand as the number of bikes *ordered* throughout the season—that is, without correcting this number for exchanges, order cancellations, or returns. The company assumed that orders were canceled owing to a combination of slow deliveries and impatient customers. Hence the number of orders placed seemed to be the best proxy for demand.

A substantial share of the company's bikes was sold out in each season. To compensate for this trend, every product manager provided an estimate of lost sales for all bikes. The average of these per-bike estimates was added to the number of orders and then actual season demand was redefined as the number of orders placed (still without correcting for exchanges, cancellations, or returns) *plus* the average estimate of lost sales. Hereafter, it is this conception of season demand that we use.

3.2. The Forecasting Team

The firm's new demand forecast would be based on the mean of the team members' forecasts as opposed to the team's consensus demand forecast before (research has shown that arithmetic means usually perform better). The team's composition would therefore be key to the accuracy of any new demand forecasting process. It was clear that all team forecasters should be well informed about the company's strategy, products, markets, competitors, and customers. To ensure that the company's overarching views were reflected in the forecast, all key functions were represented on the team. This approach is meant to result in a heterogeneous expert team featuring various perspectives and areas of specialization; it should also ensure that the experts' interests are well balanced.

It did not take much effort for Mr. Arnold to identify suitable experts for his forecasting team. He selected six of them: in addition to himself, the company's CFO, head of production and logistics, head of marketing and sales, head of R&D, and product managers. Mr. Arnold knew his team well. They had worked together on a daily basis and, more importantly, he trusted their judgment.

3.3. The Annual Demand Forecasting Meeting

Once all these hand-picked experts had arrived at Mr. Arnold's home office, the meeting could start. As always, the forecasting task was preceded by a detailed

review and discussion of the current sales season. Special attention was paid to bikes that significantly deviated from the sales plan with the aim of devising countermeasures for the next season. Next, the head of sales and business development described economic outlooks in the company's target markets as a means of assessing their growth potential. Finally, the product managers presented results from a customer survey concerning brand recognition, buying intentions, and customer-perceived competition.

After this characterization of the upcoming season was clear to all expert forecasters, the product manager for MTBs introduced next season's line-up. He began with a general overview of the MTB collection while identifying the major changes and advancements since the previous season. The product manager next presented the first bike in detail. He explained the bike's technical properties (e.g., weight, frame shape, bike handling), its components (suspension forks, breaks, gearing), the predecessor's sales curve and availability (stocking information), and the competition (benchmark with direct comparisons to competitors), as well as information regarding brand recognition and external publicity (e.g., test results, sponsoring, publications, feedback from social networks); see Appendix H. From this moment onward, all presenters were required to use a standardized presentation form and were not allowed to exert any influence on the decisions of other experts. The focal bike was then discussed by group members, and the product manager answered questions until all expert forecasters were comfortable predicting next season's demand for that bike.

Each forecaster was asked to give a demand estimate for each bike, and they were officially prohibited from sharing forecasts with other team members. The relative numbers of different frame sizes produced for a bike were assumed to be fixed (and so was not

part of any forecast). This procedure was repeated for all MTBs and, on the next day, for all RBs as well. Two full days—from early morning to late night—were required to forecast demand for all the company's bikes.


On the third day, the experts met again to discuss the outcome of the two days of forecasting. The forecasters could now, for the first time, see the group's mean demand forecast for each bike. To prevent these forecasters from being biased by the opinions of others, all demand estimates had remained undisclosed during the entire forecasting process.

4. Adjusting the Experts' Demand Forecasts

Inspired by the famous case study *Sport Obermeyer*, the local research team sought to reduce Canyon Bicycles' forecast error by using "judgmental" team forecasts (Raman and Hammond 1994). In the Sport Obermeyer case, the mean demand forecast of the expert panel served as the mean of the demand distribution. To account for demand uncertainty, the researchers used dispersion among forecasters (i.e., the standard deviation of demand forecasts). At Sport Obermeyer, however, the forecasters were found to be overconfident in their estimates and so their estimated demand distribution was, in general, too narrow. To counter this effect, the standard deviation of the demand forecast had to be doubled so that it would more reasonably proxy the distribution of actual demand.

When the head of production and logistics and the local research team analyzed Canyon Bicycles' past season's mean demand forecast (FC) versus actual season demand (AC), they could see that most observations were strictly above the diagonal line of the scatter plot; this pattern indicated that demand was (on average) being underestimated consistently (see Figure 1). Had

Figure 1. Relationship Between Mean Forecast and Actual Demand for the 2013 (Left) and 2014 (Right) Season



the forecasts been perfect, all observations would lie exactly on the diagonal line. Because the company's average forecast error was systematic, an unadjusted mean demand forecast could not be used.

The absolute forecast errors for some high-volume bikes were large, and those for some low-volume bikes were considerable. Thus it was clear that absolute forecast errors would not be useful; instead, relative forecast errors should be used for calibration.

But given that the forecast error was fairly constant over the years (on average), why not simply correct the demand forecast using the forecast error of the previous season? This approach was taken in the *O'Neill* case study (Cachon and Terwiesch 2013, Chap. 12), in which the mean demand forecast was corrected by the relative historic average forecast error. These authors proxied demand uncertainty by the standard deviation of the relative forecast errors multiplied by mean demand. With this methodology in mind, the local researchers were looking to reduce the forecast error in estimates of the demand for Canyon Bicycles and also for particular models of that product.

4.1. Product Segmentation

To determine last season's relative forecast error, it was necessary to calculate, for each bike b , the ratio of last season's actual demand $AC_{b,t-1}$ to last season's mean demand forecast $\overline{FC}_{b,t-1}$. Only then can the so-called A/F ratios be computed:

$$A/F \text{ ratio}_{b,t-1} = \frac{AC_{b,t-1}}{\overline{FC}_{b,t-1}}. \quad (1)$$

A perfect demand forecast would have an A/F ratio equal to 1. An A/F ratio smaller than 1 indicates that the demand forecast was too high, so excess inventory had to be sold at a discount. An A/F ratio greater than 1 indicates that the demand forecast was too low; in other words, the demand for bikes was not satisfied and more bikes could have been sold (see Table 1).

Actual realizations of demand are random, which means that A/F ratios across bikes are also random.

Yet, the average A/F ratios at Canyon Bicycles appeared *not* to be random. In particular, the average A/F ratio of the MTBs was considerably higher than that of the RBs; see Appendix I. Using past order data, the academic experts could identify significant differences in A/F ratios for the respective bike types.² To account for this systematic difference, they distinguished between the two segments as follows: MTBs were represented by their average A/F ratio $\overline{AF}_{t-1}^{MTB} = 1.244$, and RBs were likewise represented, $\overline{AF}_{t-1}^{RB} = 1.079$. The average A/F ratio across all bikes was $\overline{AF}_{t-1}^{ALL} = 1.162$.

To account for the forecast errors, the mean demand forecasts for the current season ($\overline{FC}_{b,t}$) were adjusted using the relevant segment's average A/F ratio during the previous season ($\overline{AF}_{t-1}^{SEG}$). Hence the adjusted mean demand forecast was now given by

$$FC_{b,t}^{ADJ} = \overline{AF}_{t-1}^{SEG} \times \overline{FC}_{b,t}. \quad (2)$$

The mean demand forecast of all bikes was then adjusted by the relative historic average forecast error to account for the group's demonstrated persistent underestimation of demand—although in this case that error is broken down by market segments (bike types). See Table 2.

Thus the previous season's relative forecast error was used to debias the group's mean demand forecast for the next season. However, the mean demand forecast is not in itself sufficient to identify the profit-optimal order quantity. Maximizing company profit—or a target service level—would require a probability distribution of demand for each bike.

4.2. Demand Forecast Volatility

According to Cachon and Terwiesch (2013), the randomness of demand can be proxied by the previous season's randomness of the A/F ratios—captured by the standard deviation (SD of $\overline{AF}_{t-1}^{SEG}$)—multiplied by the current season's mean demand forecast for each segment (see Table 3):

$$SD \text{ of } \overline{FC}_{b,t} = \overline{FC}_{b,t} \times SD \text{ of } \overline{AF}_{t-1}^{SEG}. \quad (3)$$

Table 1. Mean Demand Forecast vs. Actual Season Demand and A/F Ratio (2013)

Bike model	Category	Material	Price	Costs	$AC_{b,t-1}$	$\overline{FC}_{b,t-1}$	$A/F \text{ ratio}_{b,t-1}$
Grand Canyon	MTB	AL	5,100	3,000	221	173	1.277
Nerve	MTB	AL	1,360	800	625	456	1.371
Yellowstone	MTB	CF	680	400	966	840	1.150
Strive	MTB	CF	1,020	600	683	579	1.180
Speedmax	RB	CF	5,440	3,400	211	162	1.302
Roadlite	RB	CF	2,400	1,500	423	338	1.251
Ultimate	RB	AL	3,200	2,000	209	261	0.801
Endurace	RB	AL	560	350	787	820	0.960

Note. $\overline{FC}_{b,t-1}$ stands for last season's mean demand forecast; $AC_{b,t-1}$ stands for last season's actual demand; $A/F \text{ ratio}_{b,t-1}$ stands for last season's relative forecast error by the forecasting team.

Table 2. Adjustment of Season's 2014 Demand Forecast with 2013's A/F Ratio, Level 1

Bike model	Individual demand forecasts							$\overline{FC}_{b,t}$	$\overline{AF}_{t-1}^{SEG}$	$FC_{b,t}^{ADJ}$
	FC1	FC2	FC3	FC4	FC5	FC6	FC7			
Grand Canyon	240	110	90	180	250	175	190	176	1.244	220
Nerve	480	540	630	650	425	725	610	580	1.244	722
Yellowstone	1,100	1,250	750	1,200	1,300	870	1,300	1,110	1.244	1,381
Strive	980	950	1,100	690	750	700	650	831	1.244	1,035
Speedmax	110	180	200	130	90	120	110	134	1.079	145
Roadlite	380	325	350	410	390	430	400	384	1.079	414
Ultimate	290	350	310	360	290	300	285	312	1.079	337
Endurance	1,300	800	850	1,000	1,150	1,100	780	997	1.079	1,076

Note. FC1–FC7 stand for demand estimations of expert forecasters with FC1 = consensus forecast by Mr. Arnold; $\overline{FC}_{b,t}$ stands for the current season's mean demand forecast; $\overline{AF}_{t-1}^{SEG}$ stands for the previous season's average demand forecast error for the segment MTB, RB, respectively; $FC_{b,t}^{ADJ}$ stands for the adjusted mean demand forecast of a bike of the current season.

Table 3. Standard Deviation of the Mean Demand Forecast, Level 1

Bike model	Category	Material	$\overline{FC}_{b,t}$	SD of $\overline{AF}_{t-1}^{SEG}$	SD of $\overline{FC}_{b,t}$
Grand Canyon	MTB	AL	176	0.087	15
Nerve	MTB	AL	580	0.087	50
Yellowstone	MTB	CF	1,110	0.087	96
Strive	MTB	CF	831	0.087	72
Speedmax	RB	CF	134	0.207	28
Roadlite	RB	CF	384	0.207	79
Ultimate	RB	AL	312	0.207	65
Endurance	RB	AL	997	0.207	206

Note. $\overline{FC}_{b,t}$ stands for the current season's mean demand forecast; SD of $\overline{AF}_{t-1}^{SEG}$ stands for last season's segment-specific standard deviation of the A/F ratio for MTBs, RBs, respectively; SD of $\overline{FC}_{b,t}$ stands for current season's standard deviation of demand.

4.3. Service Level and Optimal Order Quantity

At this point, all the parameters needed to generate a probability distribution of the demand forecast were known. Yet neither is the distribution, in itself, enough to determine the optimal order quantity. One could apply the classical newsvendor technique and maximize profit by setting the expected overage costs equal to the expected underage costs. The overage costs C_o

of a bike left in stock at the end of a sales season were about 50% of its costs. The costs of a lost sales—that is, underage costs C_u —were the bike's margin: 41.18% for MTBs and 37.50% for RBs. The profit-optimal service level, or the critical ratio α , was then computed as the ratio $C_u/(C_o + C_u)$; the resulting values were 58.33% for MTBs and 54.55% for RBs. If we adopt the standard newsvendor logic, then the corresponding z-value scores were 0.210 and 0.114 for MTBs and RBs, respectively (where the z-value was obtained from a standard normal function table). The optimal order quantity $Q_{b,t}^*$ was thus given by

$$Q_{b,t}^* = FC_{b,t}^{ADJ} + z \times \text{SD of } \overline{FC}_{b,t}. \quad (4)$$

In Table 4, the optimal order quantities so derived for all eight bike models are shown.

5. Demand Forecast Accuracy

Canyon Bicycles was eager to verify the quality of the results obtained through the newly developed demand forecasting process. To test the reliability of those forecasts, the academic experts decided to check for whether their main underlying assumption—normality of season demand—would hold. If the same

Table 4. Service Level and Optimal Order Quantity, Level 1

Bike model	Category	Material	$FC_{b,t}^{ADJ}$	SD of $\overline{FC}_{b,t}$	Service level (%)	$Q_{b,t}^*$	$AC_{b,t}$
Grand Canyon	MTB	AL	220	15	58.3	223	206
Nerve	MTB	AL	722	50	58.3	732	714
Yellowstone	MTB	CF	1,381	96	58.3	1,402	1,429
Strive	MTB	CF	1,035	72	58.3	1,050	1,146
Speedmax	RB	CF	145	28	54.5	148	102
Roadlite	RB	CF	414	79	54.5	423	375
Ultimate	RB	AL	337	65	54.5	344	346
Endurance	RB	AL	1,076	206	54.5	1,099	1,257

Note. $FC_{b,t}^{ADJ}$ stands for the current season's adjusted mean demand forecast; SD of $\overline{FC}_{b,t}$ stands for current season's standard deviation of the demand forecast; service level of 58.3% and 54.5% correspond to a z value of approximately 0.210 and 0.114; $Q_{b,t}^*$ stands for the optimal order quantity; $AC_{b,t}$ stands for actual demand of the current sales season.

service level was set for all bikes and if the normality assumption was valid, then the share of in-stock bikes should correspond to the selected service level. Suppose, for instance, that the set service level $\alpha = 75\%$; then the actual season demand of 75% of the eight bike models (i.e., six bike models) should be satisfied and 25% of them (i.e., two bike models) should be out of stock.

Testing the reliability of segment-specific forecasts is best performed by tests of different service levels. Because there were eight bike models, an incremental testing interval of 12.5% was selected. Hence, the testing just described was repeated for service levels of 12.5%, 25%, 37.5%, 50%, 62.5%, 75%, and 87.5%. Plugging these incremental service levels into the normal distribution function was expected to show that, with each increment, demand for one more bike model could be satisfied. So for a service level of 12.5%, seven bike models should be out of stock; for a service level of 25%, six bike models should be out of stock; and so forth.

In a similar way, the model's goodness of fit can be determined by transferring the demand forecast of each bike into the standard z -value score. Doing so requires that the difference between the actual demand ($AC_{b,t}$) and the adjusted mean demand forecast ($FC_{b,t}^{ADJ}$) be divided by the standard deviation of the respective demand forecast (SD of $FC_{b,t}$):

$$z_{b,t} = \frac{AC_{b,t} - FC_{b,t}^{ADJ}}{\text{SD of } FC_{b,t}}. \quad (5)$$

Once all eight z -values are sorted in ascending order, they can be compared with the frequency intervals of the z -values of a standard normal distribution (see Appendix J). The predicted demand distributions were viewed as accurate whenever (i) a z -value fell into an interval of the segmented standard normal distribution and (ii) all subintervals were covered.³

6. Going Forward

At this point, Mr. Roman Arnold had a way to debias the team forecast of his experts and thereby better predict demand volatility. However, he was still not sure if all the relevant factors had been considered. Could he trust the two-segment forecast, or should he segment the analysis still further? Should he aggregate across all bikes? Was the mean demand forecast really the best point on which to base subsequent adjustments? Or would the consensus demand forecast work even better?

Mr. Arnold took another look at the standard normal distribution (Appendix J). What might the effect be of using one segment or even four different segments of bike models? Two important questions remained.

First, had the company's supply demand mismatch risk actually been reduced? Second, and perhaps of even greater interest, how much more profit did the company earn as compared with when it used the old forecasting process?

Appendix A.

Figure A.1. Mountain Bike of Season 2015



Figure A.2. Road Bike of Season 2015



Appendix B.

Table B.1. Bike of the Year Awards 2013–2014

Place	Year	Awarding magazine/Institute	Bike
1st	2014	MBR	Grand Canyon AL 5.9 W
1st	2014	Triathlon Award	Speedmax CF
1st	2014	Le Cycle	Ultimate CF SLX 9.0
1st	2013	MountainBIKE	Nerve CF Mountain
1st	2013	Roadbike	Ultimate CF SLX
1st	2013	Le Cycle	Ultimate CF SLX
1st	2013	Road Cycling UK	Ultimate CF SLX

Table B.2. Readers’ Choice Awards 2013–2014


Place	Year	Magazine	Award
1st	2014	Tour	Best bike brand
1st	2014	Bike	Best bike brand
1st	2013	Tour	Best bike brand
1st	2013	Bike	Best bike brand
2nd	2013	Freerider	Complete bike brand
3rd	2013	Roadbike	Innovation of the year

Table B.3. Design Awards 2013–2014

Year	Award
2014	Red dot design award Nerve AL
2013	IF product design award Ultimate CF SLX
2013	IF product design award gold Nerve CF
2013	Designpreis der Bundesrepublik Deutschland Speedmax CF
2013	German design award “Special Mention” Lux CF
2013	German design award “Special Mention” Ultimate CF SLX
2013	Red dot design award “Best of the Best” Ultimate CF SLX
2013	Red dot design award Nerve AL 29
2013	Red dot design award VCLS Post 2.0
2013	Good design award Grand Canyon CF SLX 29
2013	Good design award Nerve CF
2013	Good design award Speedmax CF
2013	Design and innovation award Nerve CF 9.0

Appendix C.

Figure C.1. Website of Canyon Bicycles



Appendix D.

Figure D.1. Plan of the New Production Facilities of Canyon Bicycles



Appendix E.

Figure E.1. Weekly Orders for All Bikes from Sales Season 2011 to 2014






Figure E.2. Different Sales Patterns for a MTB and a RB of the 2013 Season




Appendix F.

Figure F.1. Segmentation Level Approach



Appendix G.

Figure G.1. Cumulative Demand Distributions Across All Products (Left) and Across Two Seasons for One Main Product (Right)



Appendix H.

Figure H.1. Presentation Forecasting Meeting: Platform Overview

PURE CYCLING

NERVE AL + SERIE| Overview [Modified Platform]

- Modified platform – Better geometry [was criticized in past tests]/X12 axis [was criticized in past tests]/more spring deflections/better anti-squat values, more pedal kickback.
- Lighter wheels in the entire model range [was criticized in past tests]
- Potential cannibalization effects excepted regarding 29" und 650b Bikes
- Strongest competitors: Competitor 1/Competitor 2

Past season		Upcoming season	
Nerve AM 6.0	1.699,00€	Nerve AL + 6.0	1.749,00€
Nerve AM 7.0	1.999,00€	Nerve AL + 7.0	1.999,00€
Nerve AM 8.0 X	2.499,00€	Nerve AL + 8.0	2.599,00€
Nerve AM 9.0	2.899,00€		
Nerve AM 9.0 X	3.199,00€	Nerve AL + 9.0	2.999,00€
Nerve AM 9.0 SL	3.799,00€	Nerve AL + 9.0 SL	3.699,00€

Modell	Magazine	Country	Issue	Test result	Test winner
Nerve AM 7.0	Mountainbike	GER	June	Excellent – Weak points: Seat angle/wheels	Cube AMS 150 Pro and Radon Slide AM 7.0
Nerve AM 8.0 X	Big Bike	FR	January	Average – Weak points: Seat angle	Specialized Stumpjumper EVO
Nerve AM 9.0 SL	Terrengsykkel	NO	March	Position 2 of 8	Ghost AMR + Lector 7700


Figure H.2. Presentation Forecasting Meeting: Part Specification

PURE CYCLING

NERVE + AL 6.0|Product [Modified Platform]


Grand Canyon AL	Nerve AL + 6.0
Weight	12, 45 kg
Purchasing costs	n.a.
Sales price	1.699€ → 1.749,00€
Frame	M17-13 Nerve AM
Head Parts	Cane Creek full integrated 40 [for M17/M26] incl. Top parts
Rear Shock	Float CTD LV w/o spacer (No BV)—Evolution 190 mm
Fork	32 TALAS 26 CTD O/C—Evolution 150 mm/Tapered/QR15
R-DERAILLEUR	Shimano XT Shadow Plus RD-M786 SGS
F-DERAILLEUR	Shimano SLX FD-M661 Down Swing—High Direct Mount 3 × 10
SHIFTER-RIGHT	Shimano SLX SL-M670 right (No gear indicator)
SHIFTER-LEFT	Shimano SLX SL-M670-L (No gear indicator)
BRAKE-FRONT	Avid Elixir 3 blk
BRAKE-REAR	Avid Elixir 3 blk
Brake Rotor Front	Avid ASSY DB ROTOR/BOLTS G2CS 200
Brake Rotor Rear	Avid ASSY DB ROTOR/BOLTS G2CS 180
FRONT HUB/Wheelset	Mavic Crossride 15 mm/X12 quick release axle
REAR HUB	Mavic Crossride Disc
REAR HUB QR	Syntaxe X-12 135 + compl (incl. Washer, cone, 2pcs, o-ring, 0 mm Thread insert)
TIRE-FRONT	Continental Mountain King II 2.4 foldable Protection
TIRE-REAR	Continental Mountain King II 2.4 foldable Protection
Freewheel	Shimano Deore CS-HG 62 10 speed 11 – 36
CHAIN/WHEEL	Shimano FC-M590-10 (Black) 10 spd 42/32/24
CHAIN	KMC X10-93 silver
Stem	Race Face Ride60/70/90/100/110
Handlebar	Race Face Ride 710 × 20
Grips	Ergon GA 1 OEM Grip
Saddle	Selle Italia X1
Seat Post	Iridium—Kailloy SP-368 (27,2/ 30,9/31,6)
Seat Clamp	E9-09 MTB Seat clamp for 30, 9 mm seat post diameter

Figure H.3. Presentation Forecasting Meeting: Actual and Planed Production and Sales




Appendix I.

Figure I.1. Empirical Distribution of A/F Ratios for Sales Season 2013



Appendix J.

Figure J.1. Equal Probability Segmentation of a Standard Normal Distribution Function



Endnotes

¹ Alternatively, one could inflate the previous season's aggregated actual demand and then break the resulting forecast down into the relevant components by using that season's relative bike shares, i.e., a top-down forecast.

² The data set used for this case study was too small to yield reliable significant results from a mean difference test. However, the A/F ratios from the original Canyon Bicycles data set are statistically different for MTBs and RBs.

³ For more general cases that exhibit an "uneven" pattern, a standard normality test can be used to verify whether actual season demand satisfies the normality property.

References

- Cachon G, Terwiesch C (2013) *Matching Supply with Demand: An Introduction to Operations Management*, 3 ed. (McGraw-Hill/Irwin, Boston).
- Raman A, Hammond JH (1994) Sport Obermeyer Ltd. Harvard Business School Case 695-022, October 1994 (Revised August 2006).