The Impact of Decision Types on Revenue Management Decisions: An Experimental Study

Ayşe Kocabıyıkoğlu[†]
School of Management, Sabanci University, Istanbul 34956, Turkey,
e-mail: kocabiyikoglua@sabanciuniv.edu

Celile Itır Göğüş
Faculty of Business Administration, Bilkent University, Ankara 06800, Turkey,
e-mail: itirgogus@bilkent.edu.tr

Mert Hakan Hekimoğlu
Lally School of Management, Rensselaer Polytechnic Institute, Troy, NY 12180, USA,
e-mail: hekimm@rpi.edu

ABSTRACT

In the standard two-class revenue management model, the decision maker allocates a fixed resource between two customer classes with hierarchical prices and uncertain demand. The normative (i.e., expected revenue-maximizing) allocation is given by Littlewood's Rule, but little is known about how decision makers actually form these decisions. We report results of an experimental study that investigates revenue management decision-making. We find that subjects' behavior is influenced by the decision type. In particular, our subjects reserve more units for the high-end segment when they are asked to set the protection level (the number of units to set aside for the higher priced class) compared to when they set the booking limit (the number of units available for the lower priced class). We propose that this behavioral pattern can be explained by our subjects' different valuations of revenues from the high- and low-end sales. We also observe that when there is a change in segment prices, although decision makers adjust allocations in the direction suggested by normative theory, the magnitude of adjustments is greater (and hence closer to the normative level) when the source of the price change matches the class whose allocation they determine. [Submitted: April 19, 2016. Revised: January 1, 2017. Accepted: January 25, 2017.]

Subject Areas: Behavioral Operations Management, and Revenue Management.

INTRODUCTION

Revenue management applications in areas such as transportation and accommodation mostly rely on decision support systems, because capacity allocation and

_____

[†]Corresponding author.

225

pricing policies are implemented over many instances. However, in other industries where there are a few large events per year, such as sports and entertainment event management, there is room for individual decision making. For example, "tour operator revenue management is a largely manual operation, with analysts trawling through text reports to determine what capacity is selling well, and what is not, and manually implementing price changes or barring holidays from sale in response to this information" (Jarvis, 2002). Even in industries where revenue management decisions have been traditionally automated, the need for manual intervention has reversed this trend in the past decade (Poelt, 2011). In fact, in a survey of 37 airlines, 16 of them reported less than 20% of their overall flights were managed without any user intervention to the system (Weatherford, 2009). Because managers responsible for revenue management applications are not necessarily trained in optimization methods, studying how they actually make decisions would yield valuable insights.

In this article, we present an experimental study of the two-class revenue management problem, where a decision maker allocates a fixed resource/capacity (such as seats on a flight or rooms in a hotel) between two customer segments with sequential price levels and uncertain demands, to maximize revenues. Customers who are willing to pay a higher price for the resource (henceforth referred to as class 1 customers) arrive later on the sales horizon, and the decision maker determines either the *protection level* (the number of units reserved for the higher priced segment) or the *booking limit* (the remaining units available for sale to the lower priced class, henceforth referred to as class 2). The allocation decision is made before demand from either class is realized. We focus on the two-class model because it is the base model for revenue management, and the oldest revenue management model still in practical use. It also captures the essential features of revenue management problems such as finite supply of inventory, price differentiation, a fixed sales horizon, and uncertain demand.

Our study contributes to the existing behavioral studies of revenue management models (see, e.g., Bearden, Murphy, & Rapaport, 2008; Bendoly, 2011, 2013; Kocabiyikoglu, Gogus, & Gonul, 2015; reviewed in detail below) by investigating the impact of decision type on allocation decisions, under normatively equivalent settings. In particular, we study whether the number of units set aside for class 1 customers (or, equivalently, the number of units available for class 2 customers) differ when the decision makers set the protection level versus when they determine the booking limit, under the same revenue management problem with identical capacity, demand and price parameters. Normative theory prescribes allocation decisions should not be impacted by the decision type; because capacity is fixed, when a certain number of units is protected for class 1 customers, the remaining units are available for sale to class 2, that is, one decision implies the other. However, in our study, our subjects consistently reserve more units for class 1 customers (or, equivalently, fewer units are available to class 2 customers) when they decide the protection level, compared to when they set the booking limit. Furthermore, we observe that our subjects' allocation decisions deviate consistently from normative (i.e., expected revenue maximizing) predictions. Fewer units are set aside for the higher priced class compared to the normative level when the ratio of class 2 price to class 1 price is less than ½. When this ratio is

reversed, our subjects protect significantly more units for class 1 than the normative quantity.

Research in decision theory and behavioral operations management literature suggests, contrary to the predictions of economic theory, different sources of income can be valued differently by decision makers (Thaler, 1999; Becker-Peth, Katok, & Thonemann, 2013). Two sources of income exist within the two-class revenue management framework: revenue from sales to the higher priced class, and revenue from sales to the lower priced class. We argue that our subjects' different valuation of these two revenue streams is the driver of the decision patterns observed in our study. Specifically, adopting Becker-Peth et al.'s (2013) framework to the revenue management setting, we calculate the weights our subjects assign to class 1 sales, and observe that the subjects who are tasked with determining the protection level consistently value high-end revenues more, compared to subjects who set the booking limit. Hence, they reserve more units for class 1 customers (or, equivalently, make fewer units available to class 2 customers), as capacity allocated to the higher priced class increases in subjects' valuation of revenues from this segment.

Our experimental design also allows us to study the impact of changes in the selling prices of the two customer segments on allocation decisions. Revenue management problem environments are dynamic, with preferences, valuations and other factors changing over time (Lieberman, 2011). Furthermore, revenue management and pricing decisions are usually separated within organizations (Westermann & Lancaster, 2011), so revenue managers react to changes in prices set by the marketing department. In our study, consistent with theoretical predictions, our subjects increase the allocation to each customer class in its own selling price and decrease in the price of the other class. However, because the source of income matters to our subjects, that is, they value revenues from class 1 and class 2 sales differently, the observed adjustment in allocation levels in response to a price change is lower compared to the normative amount. Furthermore, we observe that subjects adjust allocations in greater magnitudes (hence bringing them closer to the normative levels) when the source of the price change (i.e., whether the change is in class 1 or class 2 price) and the decision type (i.e., whether the subjects set the protection level or the booking limit) match.

Related Literature

The seminal paper in the literature studying operations management models from a behavioral perspective is due to Schweitzer and Cachon (2000), who study the newsvendor problem in an experimental setting, and show that decision makers tend to choose order quantities between the normative level and the mean demand (*the pull-to-center effect*). This finding is confirmed by Benzion, Cohen, Peled, and Shavit (2008), Bolton and Katok (2008), and Bostian, Holt, and Smith (2008). The pull-to-center effect is a robust pattern under different feedback frequencies (Lurie & Swaminathan, 2009), when decisions are made in groups (Gavirneni & Xia, 2009), and across different sample types (Bolton, Ockenfels, & Thonemann, 2012). Schweitzer and Cachon (2000), Bolton and Katok (2008), and Bostian et al. (2008) argue that the pull-to-center effect can be explained by two special

cases of anchoring-and-adjustment heuristic (Kahneman & Tversky, 1979): the tendency of decision makers to anchor on mean demand and adjust insufficiently toward normative order quantities (mean anchoring heuristic), and the adjustment of current order quantities in line with previous period's demand realizations (demand chasing). Su (2008) proposes random optimization errors as the driver of the pull-to-center effect; however, Kremer, Minner, and Van Wassenhove (2010) show that random errors do not fully explain this pattern, and the decision bias is context dependent. Other explanations for the pull-to-center effect proposed in the literature include overconfidence (Ren & Croson, 2013), psychological costs of underage and overage (Ho, Lim, & Cui, 2010), impulse balance equilibria (Ockenfels & Selten, 2014), and prospect theory (Long & Nasiry, 2015).

Other experimental work on the newsvendor problem extends Schweitzer and Cachon's (2000) framework to study impact of information (Gavirneni & Isen, 2010), framing of objectives (Schultz, McClain, Robinson, & Thomas 2007), the relationship between cognitive reflection and newsvendor decisions (Moritz, Hill, & Donohue, 2013), loss aversion and source of income (Becker-Peth et al., 2013), and competition (Ovchinnikov, Moritz, & Quiroga, 2015).

There exist fewer studies of revenue management models from a behavioral perspective, compared to the newsvendor problem. Kocabiyikoglu et al. (2015) compare decision behavior under the newsvendor and the two-class revenue management models. They observe that, under mathematically equivalent settings, revenue management allocation decisions are consistently higher compared to newsvendor order quantities. They propose that this difference in behavioral patterns can be explained by the higher sensitivity of newsvendor decision makers to losses due to overage, because they incur actual costs when there are unsold units, whereas in the revenue management setting, earnings are measured in terms of revenues. Bearden et al. (2008) provide an experimental study of decision-making in revenue management, but they focus on a problem distinctly different from ours. They study behavior in a setting where subjects dynamically sell a fixed resource over a multiperiod horizon; each customer has a different valuation for the product (while we focus on two customer segments) and the firm can accept or reject a bid (unlike our problem, where demand is refused only if there are no remaining units for sale). Bendoly (2011) uses a set up similar to Bearden et al.'s (2008) to investigate the influence of arousal and stress, and forms of feedback (2013) on revenue management decisions.

As detailed above, one of the key observations in our study is that decision type impacts allocation behavior when decision makers face identical revenue management problems, with same capacity, demand and price parameters. Similar studies in behavioral operations management literature, which investigate decision-making under normatively equivalent but behaviorally different settings, include Ho and Zhang (2008), who observe that presenting the fixed fee in a pricing contract as a discount or a two-part tariff significantly impacts subjects' decisions, Katok and Wu (2009), who show that subjects behave differently under mathematically equivalent buyback and revenue sharing contracts, and Wu and Chen (2014), who propose the observed differences in behavior between two contract types are due to the way incentives are framed. In a recent study, Tokar, Aloysius, Waller, and Hawkins (2015) examine multiple types of framing effects in

inventory control decisions, and propose ways framing can be used as a debiasing technique.

Additional work in the behavioral operations management literature include studies on the supply chain model (Katok, Thomas, & Davis, 2008; Loch & Wu, 2008a; Narayanan & Moritz, 2015; Becker-Peth & Thonemann, 2016), particularly the causes of the bullwhip effect, that is, the large variations in orders as a result of small changes in demand (Croson & Donohue, 2002, 2003, 2006; Bloomfield, Gino, & Kulp, 2007). For an overview of the behavioral operations management literature and research directions, the reader is referred to Bendoly, Donohue, and Schultz (2006), Gino and Pisano (2008), Loch and Wu (2008b), and Bendoly, Croson, Goncalves, and Schultz (2010).

Structure

The rest of the article is organized as follows. The next section outlines the revenue management problem, and the details of our experimental design and laboratory implementation. "Results" section presents our experimental results. "Summary of the Results and Concluding Remarks" section provides a summary and discussion of key results, and some concluding remarks.

THE REVENUE MANAGEMENT PROBLEM AND EXPERIMENTAL DESIGN

The Revenue Management Problem

In the standard two-class revenue management model (Littlewood, 1972; Belobaba, 1989; Brumelle & McGill, 1993), the decision maker determines the allocation of a fixed resource, C, between two customer segments (class 1 and class 2) with price levels r_1 and r_2 ($r_1 > r_2$ without loss of generality). The uncertain demand from the two customer segments is given by D_1 and D_2; class 2 demand arrives before class 1 demand. The allocation decision is made before demand from either class is realized. The decision maker can either set the *protection level* (x), the portion of capacity reserved for the higher priced class, whose demand will arrive later, or the *booking limit* (b), the units available for sale to class 2. Note that determining the protection level and determining the booking limit are normatively equivalent; because capacity is fixed, one decision implies the other, that is, $b = C - x$. The decision maker's objective is revenue maximization; in typical revenue management settings, fixed costs are high and variable costs are negligible, hence, earnings are stated in terms of revenues, not profits. The expected revenues (stated in terms of the protection level) are given by

$$R(x) = r_2 \mathrm{E}\left[\min\left(D_2, C - x\right)\right] + r_1 \mathrm{E}\left[\min\left(D_1, \max\left(C - D_2, x\right)\right)\right]. \quad (1)$$

The first term in (1) gives the revenue from the low-end segment. The number of units sold to this class is constrained by its demand D_2, and the booking limit $b = C - x$. The second term is the revenue from class 1. The number of units available to this class is given by $\max(C - D_2, x)$, because the availability is uncertain, and may exceed the protection level x, if class 2 demand does not exceed the booking limit. Remark that the revenue function can equivalently be

written in terms of the booking limit as, $R(b) = r_2 \, E[\min(D_2, b)] + r_1 \, E[\min(D_1, C - \min(D_2, b))]$.

The unique normative protection level x^* solves (see, e.g., Littlewood, 1972):

$$x^* = F^{-1}\left((r_1 - r_2)/r_1\right), \tag{2}$$

where F^{-1} denotes the inverse of the cumulative distribution function of the class 1 demand D_1. Also note, the normative booking limit is given by $b^* = C - x^*$.

From (2), the normative allocation to each class increases in its own price, and decreases in the other class' price. In particular, the normative protection level increases when there is an increase in the class 1 price, and decreases when there is an increase in the class 2 price. The opposite relationships hold between the selling prices and the normative booking limit, that is, the normative booking limit increases in the class 2 price and decreases in the class 1 price.

Finally, it should be noted that, in the two-class revenue management framework, the higher priced segment is more valuable to the decision maker than class 2 sales; the low-end segment is offered on a limited basis to compete on capacity that would otherwise remain unsold. Over allocating to class 1 leaves the decision maker with unused capacity, which brings no revenue (because of the sequence of demand arrivals, units that remain unsold after class 1 demand is met cannot be offered to class 2 customers again). Under allocating to class 1 results in lost class 1 sales; in this case, units that could have been sold to the higher priced class were sold to class 2 customers earlier on the sales horizon, at a lower price.

Experimental Design

Our study employs the two-class revenue management setting, discussed above. Specifically, our subjects were asked to allocate a fixed number of airline seats between two customer segments over 30 experimental rounds, with the objective of maximizing revenues. This setting was chosen because the airline industry is the major user of revenue management practices.

We employed a $2 \times 2 \times 2$ design in our experiments. The factors distinguishing the conditions were: (1) the decision type (protection level or booking limit), (2) the price varied during the experimental rounds (class 1 or class 2), and (3) the ratio of class 2 price to class 1 price ($r_2/r_1 = \frac{1}{4}$ or $r_2/r_1 = \frac{3}{4}$). The first two factors were manipulated by a between-subject design, whereas the last factor was within-subject.

The subjects assigned to conditions *P1* and *P2* were asked to decide the *protection level* (the number of units to set aside for class 1). In the *P1* condition, we kept the same class 2 price ($r_2 = 90$) over the 30 experimental rounds, and varied the class 1 price after 15 rounds. In the *low* price treatment, we had $r_1 = 120$ (hence, $r_2/r_1 = \frac{3}{4}$); in the *high* price treatment, we had $r_1 = 360$ (hence, $r_2/r_1 = \frac{1}{4}$). In the *P2* condition, on the other hand, class 1 price remained constant for the entire experimental session, whereas class 2 price varied. We had $r_1 = 120$ and $r_2 = 90$ under the *low* price treatment, and $r_1 = 120$ and $r_2 = 30$ under the *high* price treatment.

The price treatments are denoted as *high* and *low* because the ratio of selling prices indicates the relative contribution of class 1 sales to revenues compared to

class 2 sales. For example, when $r_2/r_1 = \frac{1}{4}$, the contribution of class 1 sales is much higher than class 2 sales, compared to when $r_2/r_1 = \frac{3}{4}$, because the difference in segment prices is greater. Also remark that overallocating to class 1 (i.e., being left with unused capacity) is more detrimental to earnings under the *low* price treatment, whereas underallocating to class 1 (i.e., experiencing lost class 1 sales) hurts earnings more under the *high* price treatment.

Each price treatment lasted 15 rounds, and to check for order effects, some subjects faced the *low* price treatment first (in rounds 1–15) and the *high* price treatment in rounds 16–30, whereas the order was reversed for the rest of the subjects (order effects refer to the possibility that subjects' experiences in an experiment might bias their decisions in following experiments, Camerer, 2003).

The other two conditions in our study, *B1* and *B2*, are analogues of *P1*, and respectively *P2*, in terms of price levels, but the subjects set the *booking limit*, instead of the protection level. Specifically, in the *B1* condition, prices were set at $r_1 = 120$ and $r_2 = 90$ under the *low* price treatment, and at $r_1 = 360$ and $r_2 = 90$ under the *high* price treatment, that is, the change in prices was manipulated through a change in the class 1 price. In the *B2* condition, we kept the same class 1 price ($r_1 = 120$) throughout the experimental rounds and varied the class 2 price ($r_2 = 90$ under the *low* price treatment, and $r_2 = 30$ under the *high* price treatment). Throughout the article, we refer to price treatments under a condition *Y* as *Ylow* and *Yhigh*, with $Y = \{P1, P2, B1, B2\}$. For example, *B1low* refers to the treatment where subjects set the booking limit under price levels $r_1 = 120$ and $r_2 = 90$.

The flight for which the subjects were required to make allocation decisions had a capacity of 120 seats. Demand from both customer segments were modeled by *Uniform* distributions. Class 1 demand was distributed *Uniform*(0,80) and class 2 demand was distributed *Uniform*(50,170). These parameters were chosen because revenue management is most relevant when capacity is binding, yet ample enough to serve both segments. Capacity and demand distributions remained constant throughout the experimental rounds.

Remark that from (2), when class 1 demand is distributed *Uniform*(l,u), the normative protection level solves $x^* = u - (r_2/r_1)(u - l)$. This implies a normative protection level of $x^* = 60$ under the price treatment *high*, and of $x^* = 20$ under the price treatment *low*. Because the normative booking limit is equal to capacity minus the normative protection level ($b^* = C - x^*$), the corresponding booking limits are $b^* = 120 - 60 = 60$ under the *high* price treatment, and $b^* = 120 - 20 = 100$ under the *low* price treatment.

A total of 162 subjects participated in our study; each subject participated in exactly one condition. A summary of the experimental conditions and the number of subjects in each condition are given in Figure 1.

Laboratory implementation

All experimental sessions followed the same protocol. When the subjects arrived at the lab, they were given instructions that described the airline revenue management problem, and the experimental protocol. After all subjects read the instructions on their own, the experimenter read them aloud, using PowerPoint slides to illustrate concepts and examples, and answered questions. At the end of this training period,

Figure 1: Summary of experimental conditions.

Note: Class 1 demand ~ Uniform(0,80), Class 2 demand ~ Uniform(50,170), Capacity = 120

a short test was administered to ensure the subjects understood the instructions. Only the participants who completed the test successfully began the experimental rounds.

The experiments were conducted with specifically programmed software based on the Visual Basic for Applications (VBA) platform. At the beginning of each round, subjects were provided with demand distribution, price and capacity information for that round, and were asked to determine the protection level if they were assigned to the conditions *P1* or *P2*, and the booking limit if they were assigned to the conditions *B1* or *B2*. After making their decision, subjects learned the actual class 1 and class 2 demands for that round, and their realized revenue (total and from each class). Class 1 and class 2 demand vectors were randomly determined prior to the experiment, and were the same for every subject. In table format, the subjects could also view the number of seats sold to each class, the number of customers turned away from each class, and lost revenue due to over- and underallocating. An information message was displayed on the screen when the price parameters were changed after the 15th experimental round.

The initial five rounds were dubbed the "pilot rounds" and were used to familiarize the subjects with the interface and the task. Hence, data from this part of the experiment was not used in the subsequent analyses. Next, the subjects completed 30 experimental rounds, after which they were briefed and dismissed. Subjects were not informed of the total number of rounds prior to the experiment. None of the subjects received information about others' performance and subjects were not allowed to communicate during the experiment. Each session lasted approximately 90 minutes, including the training period and pilot rounds.

RESULTS

In this section, we first present an analysis of our subjects' allocation decisions. Then, we explore the drivers of the observed behavioral patterns.

Table 1: Descriptive statistics.

|	Average Protection Level	Average Number of Units Available for Class 2	Standard Error	Normative Protection Level
P1high	43.14	76.86	9.04	60
P1low	37.87	82.12	6.50	20
P2high	41.39	78.61	6.11	60
P2low	39.61	80.39	5.30	20
	Average Booking Limit	Average Number of Units Reserved for Class 1	Standard Error	Normative Booking Limit
B1high	84.36	35.64	4.89	60
B1low	88.13	31.86	5.89	100
B2high	81.31	38.68	6.81	60
B2low	90.22	29.78	6.26	100

Note: Letting x and b denote the protection level and booking limit decisions respectively, the number of units reserved for class 1 is equal to $C - b$, and the number of units available to class 2 is equal to $C - x$, with $C = 120$.

Figure 2: Subjects' average decisions: *P1* and *B1*.
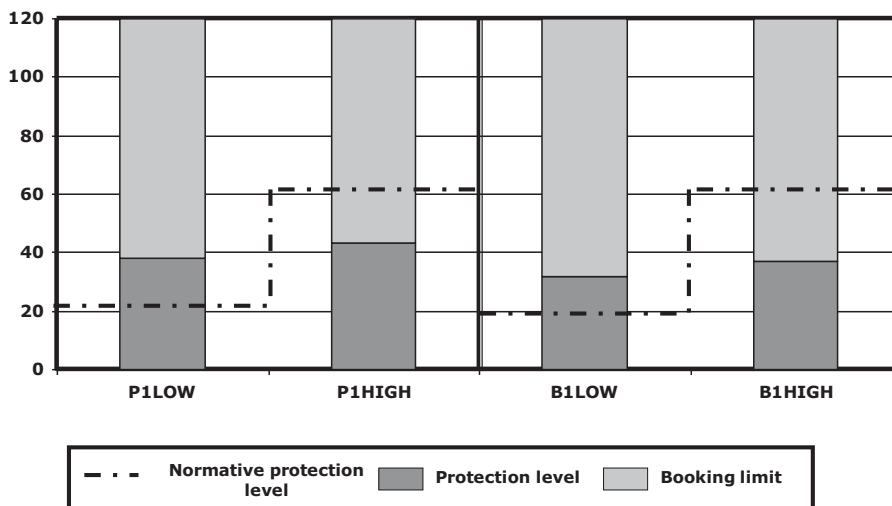
Finally, we investigate the impact of changes in prices on revenue management decisions.

Table 1 provides an overview of our subjects' average allocation decisions, as well as standard errors and normative benchmarks. Additional descriptive statistics on the overall and average decisions, including the median, maximum, and minimum, are given in Table A1. Figures 2 and 3 provide a visual summary of

Figure 3: Subjects' average decisions: *P2* and *B2*.
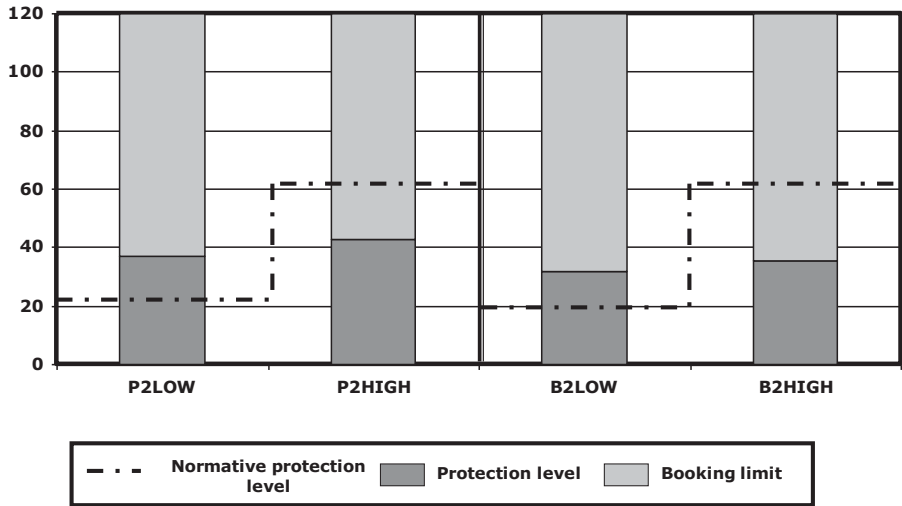
our subjects' average decisions under conditions *P1* and *B1*, respectively *P2* and *B2* (Figures A1 and A2 give the average decision per experimental round for each treatment and Figures A3 and A4 provide histograms of our subjects' decisions).

For the majority of our analyses, we used the Wilcoxon test (Siegel, 1956) for the comparisons. We checked our data for any order effects, that is, whether the order of price treatments impacted our subjects' decisions, and found none.

Next, we compare our subjects' decisions with the normative benchmarks, and investigate whether being tasked with setting the booking limit or the protection level influences actual allocation behavior. We also discuss whether the pull-to-center effect, one of the most robust behavioral patterns from the experimental newsvendor literature (see e.g., Schweitzer & Cachon, 2000; Bolton & Katok, 2008; Bostian et al., 2008), is observed in our study.

Allocation Decisions

In our study, conditions *P1* and *B1* are equivalent, except for the decision type. Specifically, our subjects faced the same revenue management problem, that is, the capacity, demand distribution and price parameters were the same under the two conditions (see Figure 1), but they determined the protection level under condition *P1*, and the booking limit under condition *B1*. The same equivalence is valid for conditions *P2* and *B2*. As noted in the previous section, normative theory predicts that the number of units allocated to each class will not be impacted by the decision type, that is, a normative protection level of x^* implies a normative booking limit of $b^* = C - x^*$. Next, we provide a comparison of the number of units reserved for class 1 customers under corresponding treatments in *P1* and *B1*, and under *P2* and *B2*, to investigate whether the behavior observed in our study mirrors this normative equivalence. It should be noted that, in our analysis, we compared, for example $x_{P1low} = 37.87$ with $C - b_{P1low} = 31.86$ (see Table 1). We could have

alternatively compared the number of units available for class 2 (e.g., b_{P1low} = 88.13 with $C - x_{P1low}$ = 82.12), and the analysis would have been the same.

A visual inspection of Figures 2 and 3 suggest that our subjects protected more units for class 1 (or equivalently, fewer units were available for class 2) when they set the protection level compared to when they determined the booking limit, under both price treatments. Wilcoxon tests confirmed these observations. In particular, pairwise tests revealed that on average, our subjects set aside significantly more units for class 1 under the *P1high* treatment than the *B1high* treatment and under the *P1low* treatment than the *B1low* treatment ($p < .001$ for both comparisons). The same trend was observed when we compared decisions under *P2* and *B2*. Our subjects protected significantly more units for class 1 under the *P2high* treatment than the *B2high* treatment, and under the *P2low* treatment than the *B2low* treatment ($p < .001$ for both comparisons).

We also compared protection level decisions under *P1* and *P2* with the normative quantities $x^* = 60$ under the *high* price treatment, and $x^* = 20$ under the *low* price treatment, and average booking limit decisions under *B1* and *B2* with $b^* = 60$ under the price treatment *high*, and $b^* = 100$ under the price treatment *low*. We found that our subjects set protection levels significantly lower than normative levels under the *P1high* and *P2high* treatments, and protection levels significantly higher than normative quantities under the *P1low* and *P2low* treatments ($p < .001$ for both comparisons). Similarly, the average booking limit decisions were significantly higher than the normative booking limit 60 under the *B1high* and *B2high* treatments, and significantly lower than the normative booking limit 100 under the *B1low* and *B2low* treatments ($p < .001$ for both comparisons). The observed behavior with respect to normative benchmarks was consistent with previous findings in the behavioral operations management literature (see Kocabiyikoglu et al., 2015).

Pull-to-center effect

In the newsvendor setting, the pull-to-center effect refers to the tendency of decision makers to choose order levels between the mean demand and the normative order quantity. Lau, Hasija, and Bearden (2014) designate this region as *the pull-to-center zone*, and argue that, if the pull-to-center effect exists in newsvendor data, "allowing for random errors in ordering decisions, the majority of orders should lie in the pull-to-center zone." Adopting their definition of the pull-to-center zone to the revenue management framework, for conditions *P1* and *P2*, we calculated the percentage of decisions that fell between the mean class 1 demand, $E[D_1] = 40$, and the normative protection level, $x^* = 60$ for the *high* price treatment, and between $E[D_1] = 40$ and $x^* = 20$ for the *low* price treatment. We observed that under the *low* price treatment, 57.67% of decisions in *P1*, and 55.61% of decisions in *P2* were in this pull-to-center zone. Under the *high* price treatment, the percentage of decisions between the mean class 1 demand and the normative protection level 60 was 50.70% in *P1*, and 57.54% in *P2*. Hence, we conclude that the pull-to-center effect was observed in our study when our subjects were tasked with setting the protection level.

To check whether the pull-to-center effect is a robust phenomenon across decision types, for conditions *B1* and *B2,* we calculated the percentage of times the

number of units reserved for the high-end segment (i.e., $x = C - b$) was between the mean class 1 demand $E[D_1] = 40$, and the normative quantities ($x^* = 120 - b^* = 60$ under the *high* price treatment, and $x^* = 120 - b^* = 20$ under the *low* price treatment). Note that we could have defined the pull-to-center zone as the region between the normative booking limit ($b^* = 100$ under price treatment *low*, and $b^* = 60$ under price treatment *high*) and the expected capacity available to class 2 customers (i.e., $C - E[D_1] = 80$), and the analysis would have been the same. We observed that under the *low* price treatment, the majority of individual decisions were in the pull-to-center zone (66.67% of decisions in *B1low* and 61.73% of decisions in *B2low*). However, only 35.56% of all decisions were between the mean class 1 demand and the normative protection level under the *B1high* treatment; the corresponding figure was 31.59% under the *B2high* treatment, suggesting that, in our study, the pull-to-center effect was not observed under all decision types, and it was not a robust pattern in the two-class revenue management setting.

The Valuation of Revenues

The analysis presented in the previous section indicates our subjects reserved significantly more units for the higher priced segment when they set the protection level, compared to when they were tasked with determining the booking limit, under equivalent revenue management settings. Furthermore, allocation decisions deviated consistently from normative benchmarks. In this section, we investigate the drivers behind these observed behavioral patterns.

The decision maker in the two-class revenue management setting has two sources of income: revenues from the high-end sales, and revenues from the low-end sales. Although economic theory posits decision makers are indifferent to the source of income (see e.g., Becker-Peth et al., 2013), there exists research that argues different sources of income can be associated with different valuations (O'Curry, 1997; Thaler, 1999; Kooreman, 2000). For example, "people value and use the winnings of an office football pool differently than they do income tax refunds" (Thaler, 1999). In the behavioral operations management literature, Becker-Peth et al. (2013) present a framework based on *mental accounting* (Kahneman & Tversky, 1979; Thaler, 1985, 1999), and propose that the source of income matters in the newsvendor setting. In particular, they distinguish between income from sales and income from returns to the supplier, and show that their subjects' valuation of earnings from these two sources are different.

In this section, we argue that, in our study, decision makers value revenues from class 1 and class 2 sales differently, and we propose this pattern as the driver of the different allocation behavior observed when our subjects set the protection level and the booking limit, and of the deviations from normative theory. To this end, adopting Becker-Peth et al.'s (2013) framework for the revenue management setting, we denote the weight our subjects assign to revenues from class 1 with β. Hence, the revenue function (1) can be written as $R(x) = r_2 E[\min(D_2, C - x)] + \beta r_1 E[\min(D_1, \max(C - D_2, x))]$. It follows that, for a given subject's protection level decision x, β solves,

Table 2: Average weight given to class 1 revenues.

	Low Price Treatment	*High* Price Treatment
P1	2.4	0.9
P2	1.98	0.78
B1	1.42	0.54
B2	1.43	0.59

$$x = F^{-1}\left((\beta r_1 - r_2)/\beta r_1\right), \tag{3}$$

where recall, F^{-1} denotes the inverse of the cumulative distribution function of the class 1 demand, and r_1 and r_2 are the prices of the two customer classes. From (3), the number of units reserved for the high-end customers increases in the valuation of revenues from this segment.

Table 2 provides the average weight our subjects assigned to revenues from class 1 for all our experimental conditions. Note that $\beta > 1$ implies that our subjects valued revenues from class 1 more compared to revenues from class 2, and vice versa. The income parameter β also captures the decision maker's perception of the unit contribution of sales from the high-end segment. In particular, $\beta > 1$ implies that the decision maker believes the unit contribution of sales from the high-end segment is more than the selling price of this class. When $\beta < 1$, on the other hand, the decision maker perceives the contribution of sales from class 1 to be lower than its actual price.

From Table 2, the weights our subjects assigned to revenues from sales to class 1 customers were, on average, significantly higher than one under the *low* price treatment, for all experimental conditions in our study ($p < .001$ for all comparisons). Under the *high* price treatment, on the other hand, the observed values for β were significantly below 1 ($p < .001$), except under *P1high*, where this difference was not significant ($p > .1$).

As outlined in the previous section, our subjects set aside significantly more units for the high-end segment compared to normative benchmarks under the *low* price treatment in all experimental conditions, whereas fewer units were allocated to class 1 compared to the normative quantities under the *high* price treatment. This observation is consistent with the *too low/too high pattern* observed in the behavioral studies of the newsvendor problem (see e.g., Schweitzer & Cachon, 2000). In the newsvendor context, the too low/too high pattern refers to the tendency of decision makers to order more than the normative quantity when unsold units are more detrimental to profits, and to order fewer units compared to the expected profit maximizing quantity when turning away customers hurt earnings more. Various explanations have been proposed for this pattern in the behavioral operations management literature. For example, Ren and Croson (2013) posit that newsvendors underorder when lost sales hurt earnings more, and overorder when unsold units are more costly, because they believe the distribution of demand to be less variable than it actually is; whereas Long and Nasiry (2015) propose an explanation based on prospect theory. In our study, this tendency of decision makers to underestimate the impact of overage and underage is captured with the income

parameter β. As mentioned above, β can be interpreted as the decision makers' perception of the unit contribution of sales from the high-end segment. Under the *high* price treatment, the observed β values in our study were consistently lower than 1, which implies that, our subjects believed the unit contribution of the high-end sales to be lower than its selling price. In other words, they perceived the cost of turning away class 1 customers, that is, the unit margin $p–c$, to be lower than it actually is; hence, they set aside fewer units for the high-end segment compared to normative allocations. Under the *low* price treatment, on the other hand, our subjects consistently perceived the contribution of the high-end sales to be higher than its actual price, and underestimated the cost of over allocating to class 1, which led to more units being set aside for class 1 (or equivalently, fewer units being available to class 2), compared to the normative benchmarks.

In order to investigate whether the different valuations of revenues from the high- and low-end segments can explain the different allocation behavior observed when our subjects set the protection level and the booking limit, we compared β values under the same revenue management setting, across decision types. In particular, we compared average weights assigned to class 1 revenues in condition *P1* with the weights observed in the corresponding price treatments in *B1* (2.4 vs. 1.42 under the *low* price treatment, and 0.9 vs. 0.54 under the *high* price treatment, see Table 2). We observed that the valuation of high-end revenues was significantly higher in *P1* compared to *B1* under both price treatments ($p < .001$ for both comparisons). The same pattern was valid when we compared the valuations of high-end revenues under *P2* and *B2*. Specifically, our subjects valued revenues from class 1 sales significantly more when they set the protection level compared to when they determined the booking limit under both *low* (1.98 vs. 1.43; $p < .001$), and *high* price treatments (0.78 vs. 0.59; $p < .001$); hence, they reserved significantly more units for the high-end segment.

This observation is consistent with previous findings in decision theory and behavioral operations management that shows the valuation of sources of income might differ across decision makers. Thaler (1999) notes that considerable variations are observed among households in terms of how they value different sources of income, as well as how they group different expenditures into budgets and allocate wealth into accounts. Becker-Peth et al. (2013), in their study of newsvendor ordering behavior outlined above, observe that, at the subject level, there exists significant variations regarding the valuations of income from sales and income from returns.

In our study, it was the framing of the decisions that induced behavior in line with a mental accounting framework; in particular, the source of income mattered to our subjects. In a behavioral study of the newsvendor problem, Chen, Kok and Tong (2013) also find that the framing of the decision problem (in their study, the payment scheme) leads to order patterns consistent with mental accounting.

Impact of Price Changes

Our experimental design also allows us to investigate the impact of changes in selling prices on revenue management decisions. As discussed above, in order to maximize expected revenues, decision makers should set higher protection levels

(or lower booking limits), when faced with an increase in the class 1 price, and they should set lower protection levels, or higher booking limits, when faced with an increase in the class 2 price.

In the experimental conditions *P1* and *B1*, class 1 price is higher under the price treatment *high* compared to the price treatment *low* ($r_1 = 360$ vs. $r_1 = 120$), whereas class 2 price remains the same ($r_2 = 90$). Similarly, under *P2* and *B2*, class 2 price is higher under price treatment *low* compared to treatment *high* ($r_2 = 90$ vs. $r_2 = 30$), whereas class 1 price remains the same at $r_1 = 120$. Hence, more units should be reserved for the high-end segment under *P1high* compared to *P1low*, and under *P2high* compared to *P2low*. Similarly, normative theory predicts protection levels should be lower and booking limits should be higher under the *low* price treatment of *B1* and *B2* compared to the *high* price treatment.

Indeed, when we compared our subjects' average decisions across *high* and *low* price treatments in *P1* and *P2* (i.e., $x_{P1high} = 43.14$ with $x_{P1low} = 37.87$ and $x_{P2high} = 41.39$ with $x_{P2low} = 39.61$), we observed that our subjects set significantly higher average protection levels under the *high* price treatment compared to the *low* price treatment ($p < .001$ for both comparisons). Similarly, under *B1* and *B2*, our subjects set, on average, significantly lower booking limits under the *high* price treatment compared to the *low* price treatment ($b_{B1high} = 84.36$ vs. $b_{B1low} = 88.13$ and $b_{B2high} = 81.31$ with $b_{B2low} = 90.22$; $p < .001$ for both comparisons). Overall, our subjects behaved in line with normative theory, and increased the allocation to each customer class in its own selling price, and decreased in the price of the other class.

Although the *direction* of the change in allocation decisions with respect to changing selling prices was in line with the normative theory, the *magnitude* of change was significantly lower compared to the normative amount under all our experimental conditions. In our study, the normative protection levels under the *low*, respectively *high*, price treatments were, $x^* = 20$, respectively $x^* = 60$, whereas the corresponding booking limits were $b^* = 100$, respectively $b^* = 60$. Hence, in order to maximize expected revenues, decision makers should change their allocations by 40 units in response to a change in price treatments. However, the observed difference was significantly lower than 40 under all our experimental conditions ($p < .0001$ for all comparisons; see Figure 4 for a visual summary). As discussed in section "The Valuation of Revenues", because class 1 revenues are valued more than class 2 revenues under the *low* price treatment, protection levels are higher (or, equivalently, booking limits are lower), compared to the normative benchmarks. On the other hand, the lower valuation of high-end revenues under the *high* price treatment results in lower protection levels (or, equivalently higher booking limits) compared to the normative quantity. These two patterns drive down the adjustment in allocation decisions across price treatments; hence, observed differences are smaller than the normative difference 40.

We also observed that the magnitude of change in allocation decisions in response to a change in prices was different across experimental conditions. Although the different valuations of revenues from the two customer segments suggest the adjustment in allocations would be lower than the normative level 40, it cannot predict the differences in magnitude across experimental conditions. To pursue this

Figure 4: Average change in decisions across price treatments.
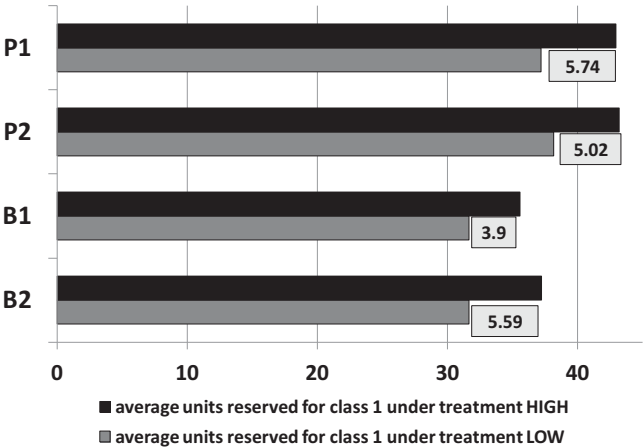
line of investigation, we first compared the amount of change in allocations across decision types. Specifically, we compared the change in decisions under *P1* with *B1*, and *P2* with *B2*. We observed that when conditions *P1* and *B1* were compared, our subjects increased the number of units reserved for class 1 significantly more in response to an increase in class 1 price when they set the protection level compared to when they determined the booking limit (5.74 vs. 3.94; $p < .001$). On the other hand, the number of units available for the low-end segment increased significantly more when the class 2 price increased under the *B2* condition compared to the *P2* condition (5.02 vs. 5.59; $p < .001$). These results suggest that there does not exist a consistent relationship between the adjustment in allocation decisions and decision type because, when conditions *P1* and *B1* were considered, subjects who set the protection level increased the number of units reserved for the high-end customers more, whereas the adjustment in allocation decisions was greater in *B2* compared to *P2*.

We also compared the change in allocation decisions within the same decision type. In particular, we compared the average adjustment in subjects' protection-level decisions under the *P1* condition with the *P2* condition, and the average adjustment in booking limit decisions when the price treatment changed under the *B1* condition with the *B2* condition. Wilcoxon tests revealed that the average adjustment in protection level decisions was significantly higher under the *P1* condition compared to the *P2* condition (5.74 vs. 5.02; $p < .001$). Similarly, our subjects increased booking limits significantly more under the *B2* condition compared to the *B1* condition when the price treatment changed from *high* to *low* (5.59 vs. 3.94; $p < .001$).

Remark that in condition *P1*, subjects determined the protection level, and the change in prices was manipulated through a change in the class 1 price ($r_1 = 360$ in the price treatment *high* and $r_1 = 120$ in the price treatment *low*). Similarly, in condition *B2*, the decision type was the booking limit, and the class 2 price increased from $r_2 = 30$ in the *high* price treatment to $r_2 = 90$ in the *low*

price treatment, whereas the class 1 price remained the same ($r_1 = 120$). The observations outlined above imply our subjects changed their allocations in greater magnitudes in response to a change in prices, under conditions where the decision type (protection level, booking limit) and the source of price change (class 1, class 2) matched.

SUMMARY OF THE RESULTS AND CONCLUDING REMARKS

In this article, we presented an experimental study of the two-class revenue management problem. The main takeaway from our study is that the decision type matters in revenue management decision-making. In particular, although normative theory prescribes the same amount of capacity should be reserved for the higher priced class (or equivalently, should be available to the lower priced class) whether the decision makers are tasked with setting the protection level or the booking limit, our subjects consistently reserved more units for the higher priced class when they were asked to set the protection level, and the number of available units for class 2 was higher when the subjects decided the booking limit. We also observed that actual allocation decisions deviated significantly from normative benchmarks.

We argue that these allocation patterns were observed, because, contrary to the predictions of standard economic theory, the source of their income matters to our subjects. In particular, subjects who were tasked with setting the protection level valued revenues from class 1 sales more, compared to those who determined the booking limit; hence, they reserved more units for the high-end segment. This finding is consistent with previous work in decision theory and behavioral operations management that show the type of decision maker or the framing of the decision problem can induce behavioral patterns consistent with mental accounting, in particular, different valuations of sources of income (Thaler, 1999; Becker-Peth et al., 2013; Chen et al., 2013). Furthermore, when the ratio of class 2 price to class 1 price was greater than ½, our subjects valued the revenues from the high-end segment more than the income from the low-end segment, or equivalently, they believed the unit contribution of sales from the high-end segment to be higher than its actual price, which led to more units being protected for class 1 customers compared to normative quantities. When the ratio of prices was reversed, revenues from class 2 sales were valued more, or the subjects believed the contribution of class 1 sales to revenues to be lower than its selling price; hence, fewer units were allocated to the high-end customers, compared to normative quantities.

Our experimental design also allows us to study the behavior of allocation decisions with respect to changes in selling prices. We observed that the direction of change in allocations was in line with the normative theory. However, because our subjects valued revenues from the high- and low-end segments differently under different price ratios, the actual amount of change was lower compared to the expected revenue maximizing quantity. Furthermore, the adjustment in decisions was higher when the source of the price change (class 1, class 2) and the decision type (protection level, booking limit) matched.

Our study contributes to the existing work in behavioral operations management literature which compares normatively equivalent but behaviorally different decision-making settings (see e.g., Ho & Zhang, 2008; Katok & Wu, 2009; Wu &

Chen, 2014; Kocabiyikoglu et al., 2015), by investigating the impact of decision type on allocation behavior under the same revenue management problem, with identical capacity, demand and price parameters. Moreover, our results, that compare actual and normative decisions, extend previous behavioral studies of revenue management models (Bearden et al., 2008; Kocabiyikoglu et al., 2015) and the newsvendor problems (e.g., Schweitzer & Cachon, 2000; Bolton & Katok, 2008; Bostian et al., 2008), by proposing the source of income, in particular, the different valuation of revenues from the high- and low-end sales, as the driver of observed allocation patterns.

Our results suggest that in order to reap full benefits of a revenue management system, companies might establish training and education programs to make revenue managers aware of the biases identified in this study, and develop decision support systems that highlight the equivalency of revenue management problems under different decision types and price changes, whether revenue management applications are completely manual or partly automated with manual interventions.

The results presented in this article can be extended in a number of ways. An interesting extension of our results might be to models with more than two customer classes; in such models, the definitions of protection level and booking limit are slightly different, and different behavioral patterns might be observed under different decision types. It might also be valuable to study changes in other problem parameters such as demand and capacity. We also believe our results regarding the impact of decision type can provide insights about decision making under other operations management models with similar settings. For example, Talluri and van Ryzin (2004) argue allocating capacity and managing prices are normatively equivalent in the revenue management context, because both limiting supply or increasing prices can achieve reducing sales. Our results suggest the decision type/frame may lead to significantly different decisions under such settings. Finally, although pricing and revenue management decisions have traditionally been treated separately, recent work has demonstrated the benefits of these two functions merging (Westermann & Lancaster, 2011; Kocabiyikoglu, Popescu, & Stefanescu, 2014); hence studying joint price and allocation decisions might yield valuable insights.

REFERENCES

Bearden, J. N., Murphy, R. O., & Rapoport, A. (2008). Decision biases in revenue management: Some behavioral evidence. *Manufacturing and Service Operations Management*, *10*(4), 625–636.

Becker-Peth, M., Katok, E., & Thonemann, U. W. (2013). Designing buyback contracts for irrational but predictable newsvendors. *Management Science*, *59*(8), 1800–1816.

Becker-Peth, M., & Thonemann, U. W. (2016). Reference points in revenue sharing contracts—How to design optimal supply chain contracts. *European Journal of Operational Research*, *249*(3), 1033–1049.

Belobaba, P. P. (1989). Application of a probabilistic decision model to airline seat inventory control. *Operations Research*, *37*, 183–197.

Bendoly, E. (2011). Linking task conditions to physiology and judgment errors in revenue management systems. *Production and Operations Management*, *20*(6), 860–876.

Bendoly, E. (2013). Real-time feedback and booking behavior in the hospitality industry: Moderating the balance between imperfect judgment and imperfect prescription. *Journal of Operations Management*, *31*, 62–71.

Bendoly, E., Croson, R., Goncalves, P., & Schultz, K. L. (2010). Bodies of knowledge for research in behavioral operations. *Production and Operations Management*, *19*(4), 434–452.

Bendoly, E., Donohue, K., & Schultz, K. L. (2006). Behavior in operations management: Assessing recent findings and revisiting old assumptions. *Journal of Operations Management*, *24*, 737–752.

Benzion, U., Cohen, Y., Peled, R., & Shavit, T. (2008). Decision-making and the newsvendor problem: An experimental study. *Journal of the Operational Research Society*, *59*, 1281–1287.

Bloomfield, R. J., Gino, F., & Kulp, S. L. (2007). *Behavioral causes of the bullwhip effect in a single echelon*. Working paper, Ithaca NY: Cornell University.

Bolton, G. E., & Katok, E. (2008). Learning-by-doing in the newsvendor problem: A laboratory investigation of the role of experience and feedback, *Manufacturing and Service Operations Management*, *10*(3), 519–538.

Bolton, G. E., Ockenfels, A., & Thonemann, U. W. (2012). Managers and students as newsvendors. *Management Science*, *58*(12), 2225–2233.

Bostian, A. A., Holt, C. A., & Smith, A. M. (2008). Newsvendor "pull-to-center" effect: Adaptive learning in a laboratory experiment. *Manufacturing and Service Operations Management*, *10*(4), 590–608.

Brumelle, S. L., & McGill, J. I. (1993). Airline seat allocation with multiple nested fare classes. *Operations Research*, *41*, 127–137.

Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton, NJ: Princeton University Press.

Chen, L., Kok, A. G., & Tong, J. D. (2013). The effect of payment schemes on inventory decisions: The role of mental accounting. *Management Science*, *59*, 436–451.

Croson, R., & Donohue, K. (2002). Experimental economics and supply chain management. *Interfaces*, *32*, 74–82.

Croson, R., & Donohue, K. (2003). The impact of POS data sharing on supply chain management: An experimental study. *Production and Operations Management*, *12*, 1–11.

Croson, R., & Donohue, K. (2006). Behavioral causes of the bullwhip effect and the observed value of inventory information. *Management Science*, *52*, 323–336.

Gavirneni, S., & Isen, A. M. (2010). Anatomy of a newsvendor decision: Observations from a verbal protocol analysis. *Production and Operations Management*, *19*(4), 453–462.

Gavirneni, S., & Xia, Y. (2009). Anchor selection and group dynamics in newsvendor decisions: A note. *Decision Analysis*, *6*(2), 87–97.

Gino, F., & Pisano, G. (2008). Toward a theory of behavioral operations. *Manufacturing and Service Operations Management*, *10*(4), 676–691.

Ho, T. H., Lim, N., & Cui, T. M. (2010). Reference dependence in multilocation newsvendor models: A structural analysis. *Management Science*, *56*(11), 1891–1910.

Ho, T. H., & Zhang, J. (2008). Designing pricing contracts for boundedly rational customers: Does the framing of the fixed fee matter? *Management Science*, *54*(4), 686–700.

Jarvis, P. (2002). Introducing yield management into a new industry. *Journal of Revenue and Pricing Management*, *1*(1), 67–75.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–292.

Katok, E., Thomas, D., & Davis, A. (2008). Inventory service-level agreements as coordination mechanisms: The effect of review periods. *Manufacturing and Service Operations Management*, *10*(4), 609–624.

Katok, E., & Wu, D. (2009). Contracting in supply chains: A laboratory investigation. *Management Science*, *55*(12), 1953–1968.

Kocabiyikoglu, A., Gogus, C. I., & Gonul, M. S. (2015). Revenue management vs. newsvendor decisions: Does behavioral response mirror normative equivalence? *Production and Operations Management*, *24*(5), 750–761.

Kocabiyikoglu, A., Popescu, I., & Stefanescu, C. (2014). Pricing and revenue management: The value of coordination. *Management Science*, *60*(3), 720–752.

Kooreman, P. (2000). The labeling effect of a child benefit system. *American Economic Review*, *90*, 571–583.

Kremer, M., Minner, S., & Van Wassenhove, L. N. (2010). Do random errors explain newsvendor behavior?. *Manufacturing and Service Operations Management*, *12*(4), 673–681.

Lau, N., Hasija, S., & Bearden, N. (2014). Newsvendor pull-to-center reconsidered. *Decision Support Systems*, *58*, 68–73.

Lieberman, W. (2011). From yield management to price optimization: Lessons learned. *Journal of Revenue and Pricing Management*, *10*(1), 40–43.

Littlewood, K. (1972). Forecasting and control of passenger bookings. *AGIFORS Symposium Proc. 12*, Nathanya, Israel.

Loch, C. H., & Wu, Y. (2008a). Social preferences and supply chain performance: An experimental study. *Management Science*, *54*(11), 1835–1849.

Loch, C. H., & Wu, Y. (2008b). *Behavioral operations management*. Boston MA: New Publishers.

Long, X., & Nasiry, J. (2015). Prospect theory explains newsvendor behavior: The role of reference points. *Management Science*, *61*(12), 3009–3012.

Lurie, N. H., & Swaminathan, J. M. (2009). Is timely information always better? The effect of feedback frequency on decision making. *Organizational Behavior and Human Decision Processes*, *108*(2), 315–329.

Moritz, B., Hill, A. V., & Donohue, K. (2013). Individual differences in the newsvendor problem: Behavior and cognitive reflection. *Journal of Operations Management*, *31*, 72–85.

Narayanan, A., & Moritz, B. B. (2015). Decision making and cognition in multi-echelon supply chains: An experimental study. *Production and Operations Management*, *24*(8), 1216–1234.

O'Curry, S. (1997). *Income source effects*. Working paper, Chicago: DePaul University.

Ockenfels, A., & Selten, R. (2014). Impulse balance in the newsvendor game. *Games and Economic Behavior*, *86*, 237–247.

Ovchinnikov, A., Moritz, B., & Quiroga, B. F. (2015). How to compete against a behavioral newsvendor. *Production and Operations Management*, *24*(11), 1783–1793.

Poelt, S. (2011). The rise and fall of RM. *Journal of Revenue and Pricing Management*, *10*(1), 23–25.

Ren, Y., & Croson, R. (2013). Overconfidence in newsvendor orders: An experimental study, *Management Science*, *59*(11), 2502–2517.

Schultz, K. L., McClain, J. O., Robinson, L. W., & Thomas, J. (2007). *The use of framing in inventory decisions*. Working paper, Ithaca NY: Cornell University.

Schweitzer, M. E., & Cachon, G. P. (2000). Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science*, *49*(3), 404–420.

Siegel, S. (1956). *Nonparametric statistics for behavioral sciences*. New York: McGraw-Hill.

Su, X. (2008). Bounded rationality in newsvendor models. *Manufacturing and Service Operations Management*, *10*(4), 566–589.

Talluri, K., & van Ryzin, G. (2004). *The theory and practice of revenue management*. Boston: Kluwer Academic Publishers.

Thaler, R. (1985). Mental accounting and consumer choice. *Marketing Science*, *4*(3), 199–214.

Thaler, R. (1999). Mental accounting matters. *Journal of Behavioral Decision Making*, *12*(3), 183–206.

Tokar, T., Aloysius, J., Waller, M., & Hawkins, D. L. (2015). Exploring framing effects in inventory control decisions: Violations of procedure invariance. *Production and Operations Management*, *25*(2), 306–329.

Weatherford, L. R. (2009). Survey results from airlines that use revenue management software systems. *Journal of Revenue and Pricing Management*, *8*(4), 323–329.

Westermann, D., & Lancaster, J. (2011). Improved pricing and integration with revenue management - The next step towards improved revenues. *Journal of Revenue and Pricing Management*, *10*(3), 199–210.

Wu, D. Y., & Chen, K. Y. (2014). Supply chain contract design: Impact of bounded rationality and individual heterogeneity. *Production and Operations Management*, *23*(2), 253–268.

APPENDIX

Table A1: Additional descriptive statistics.

	Subjects' Responses			Subjects' Average Responses		
	Min	Max	Median	Min	Max	Median
P1high	4	80	42	30	76	41.66
P1low	0	80	35	16	66.67	38.67
P2 high	5	80	45	28.73	59.93	43.67
P2 low	2	78	37	15.8	49	39.8
B1 high	50	120	85	70.4	102.33	84.6
B1 low	40	120	90	75.67	103.87	87.67
B2 high	40	120	81	61.4	104.67	84.2
B2 low	40	120	90	66.67	103.33	89.2

Note: Subjects' decisions were in terms of the protection level under conditions *P1* and *P2*, and in terms of the booking limit under conditions *B1* and *B2*.

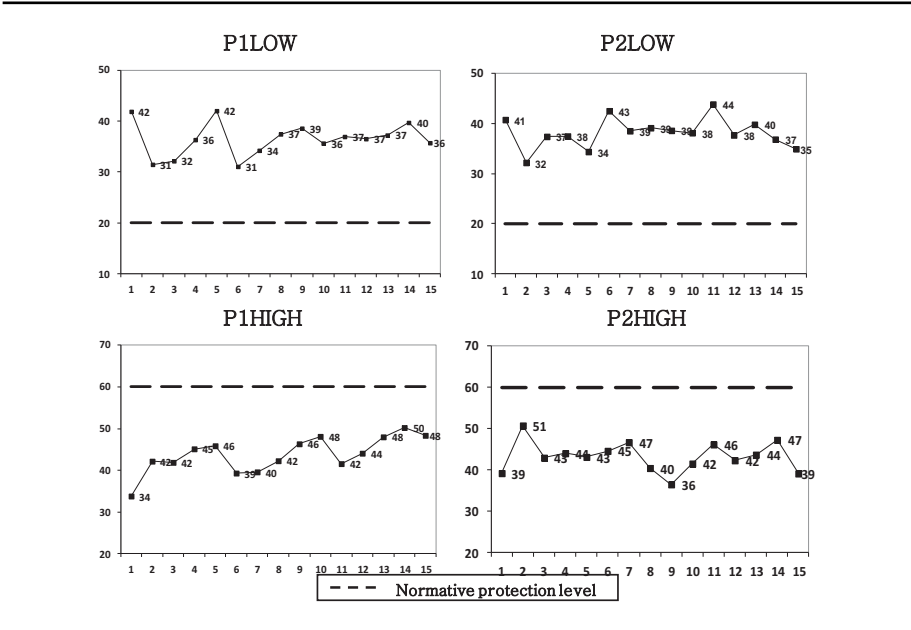**Figure A1:** Subjects' average protection level decisions per round: *P1* and *P2* conditions.

Figure A2: Subjects' average booking limit decisions per round: *B1* and *B2* conditions.
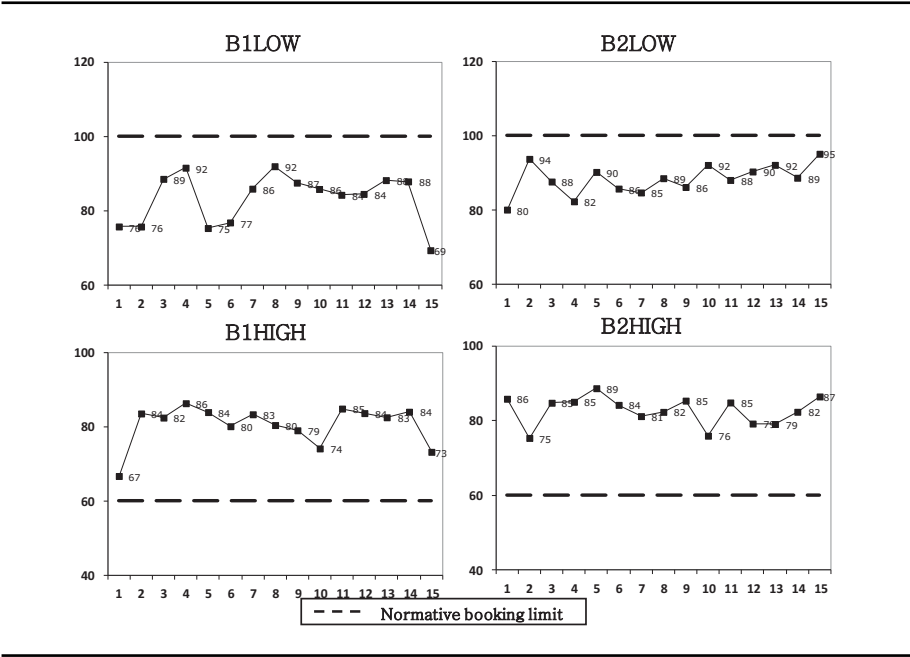

Figure A3: Histograms of subjects' protection level decisions: *P1* and *P2* conditions.
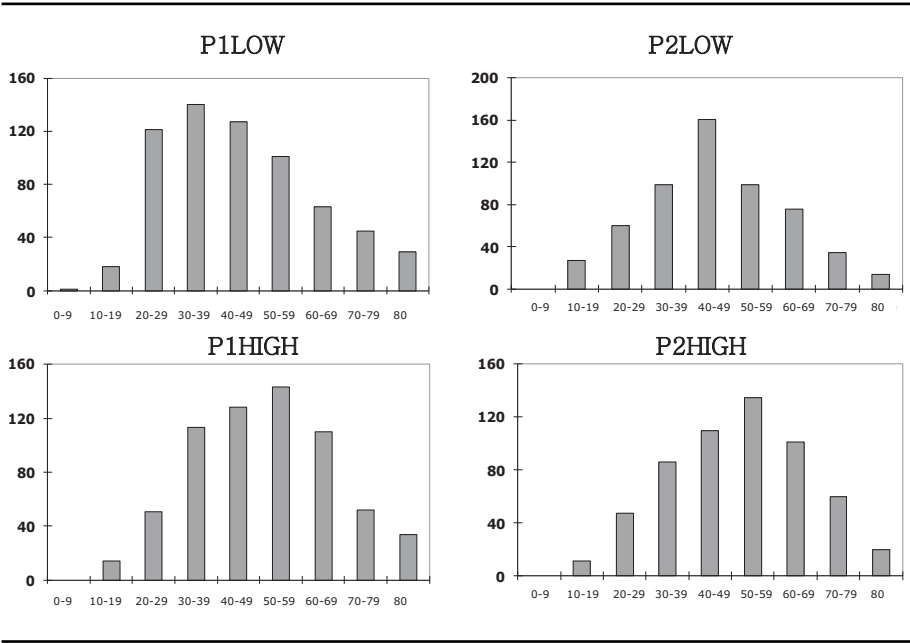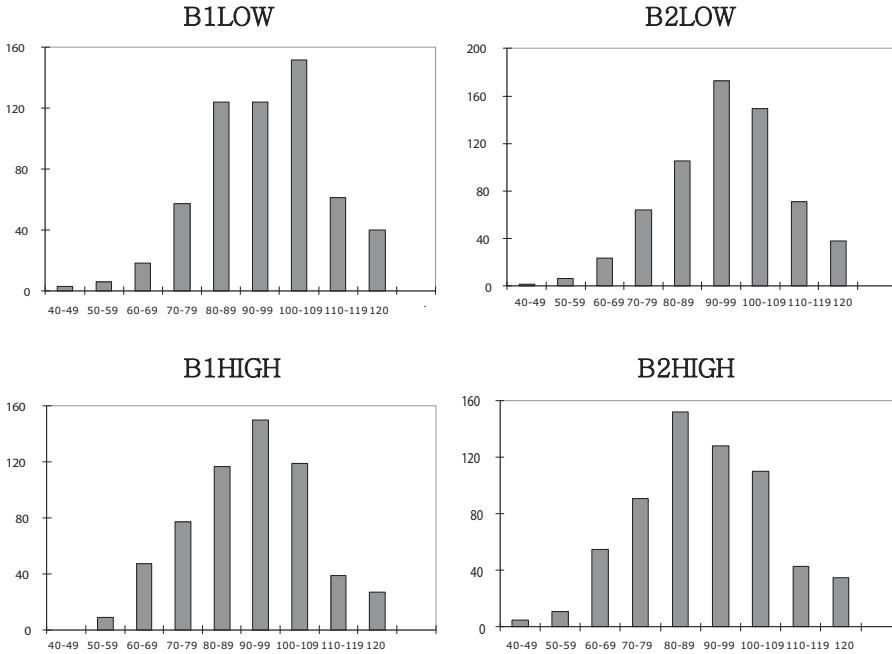
Figure A4: Histograms of subjects' booking limit decisions: *B1* and *B2* conditions.

Ayşe Kocabıyıkoğlu is an Assistant Professor of Decision Sciences at Sabancı University in Istanbul, Turkey, and holds a PhD in Decision Sciences from INSEAD, and MS and BS degrees from Middle East Technical University. Her research is at the interface of decision sciences and behavioral operations, and her research interests include behavioral operations management, pricing and revenue management, and dynamic decision making. Her work has featured in leading academic journals such as *Management Science, Operations Research* and *Production & Operations Management.*

Celile Itır Göğüş is an Assistant Professor of Management at Bilkent University in Ankara, Turkey. She has completed her PhD in Management and MS in Human Resource Management at Mays Business School, Texas A&M University and BS in Management at Bilkent University. Her research interests include behavioral applications of organizational phenomena including managerial behavior and decision making. Her research has appeared in journals such as the *Production & Operations Management*, *Journal of Applied Psychology, Journal of Business Ethics, Small Group Research, Applied Psychology* and also as book chapters.

M. Hakan Hekimoglu is an Assistant Professor of Supply Chain Management and Analytics at the Lally School of Management at Rensselaer Polytechnic

Institute. Dr. Hekimoglu holds a PhD degree in Supply Chain Management from Syracuse University. He has received his BS in Industrial Engineering and MS in Finance from Bilkent University in Ankara, Turkey, and his MA in Economics from Syracuse University. Dr. Hekimoglu's research interests include interdisciplinary problems at the interface of supply chain management, finance, and business analytics, with an emphasis on managing uncertainty and risk in complex supply chain problems.