ARTICLE IN PRESS

European Journal of Operational Research xxx (xxxx) xxx

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Decision Analysis

Risk-averse algorithmic support and inventory management

Pranadharthiharan Narayanan [a], Jeeva Somasundaram [b], Matthias Seifert [b,*]

[a] Nova School of Business and Economics, Universidade NOVA de Lisboa, Campus de Carcavelos, 2775-405 Carcavelos, Portugal
[b] IE Business School, IE University, Calle Maria de Molina 12, 5° 28006 Madrid, Spain

ARTICLE INFO

JEL codes:
D81

Keywords:
decision analysis
Algorithm
Risk aversion
Decision support systems
Anchoring

ABSTRACT

We study how managers allocate resources in response to algorithmic recommendations that are programmed with specific levels of risk aversion. Using the anchoring and adjustment heuristic, we derive our predictions and test them in a series of multi-item newsvendor experiments. We find that highly risk-averse algorithmic recommendations have a strong and persistent influence on order decisions, even after the recommendations are no longer available. Furthermore, we show that these effects are similar regardless of factors such as source of advice (i.e., human vs. algorithm) and decision autonomy (i.e., whether the algorithm is externally assigned or chosen by the subjects themselves). Finally, we disentangle the effect of risk attitude from that of anchor distance and find that subjects selectively adjust their order decisions by relying more on algorithmic advice that contrasts with their inherent risk preferences. Our findings suggest that organizations can strategically utilize risk-averse algorithmic tools to improve inventory decisions while preserving managerial autonomy.

1. Introduction

Firms increasingly rely on algorithms for inventory management tasks such as forecasting, pricing, and product replenishment (Kesavan & Kushwaha, 2020). Adoption of such algorithmic recommendations helps firms improve operational efficiency and reduce costs. For instance, in 2019, food retailers who used algorithmic inventory planning reduced stock shortages of fresh produce by 25 % and improved their gross margins by 9 % (Föbus et al., 2019). Despite the growing popularity of algorithmic decision support, recent research shows that people tend to make judgmental adjustments to algorithmic recommendations or override them entirely, even if this comes at the cost of lower performance (Dietvorst et al., 2015; Fildes et al., 2009; Khosrowabadi et al., 2022). Studying this behavioral tendency and finding ways to overcome what is now dubbed as "algorithm aversion" has become an important research agenda in the behavioral operations community (see, for instance, Lin et al., 2023). The present paper extends this line of research by studying behavioral responses to algorithms that are calibrated with specific levels of risk aversion in the context of inventory management decisions.

More generally, risk preferences play a fundamental role in inventory management (Chen et al., 2007). Importantly, risk-averse inventory decisions help reduce volatility in profits. For instance, in many scenarios which involve significant inventory holding costs (e.g., time-sensitive medical supplies), short product lifecycles (e.g., consumer electronics), high demand uncertainty (e.g., fashion retail), capital constraints (e.g., small businesses) or perishable goods with no salvage value (e.g., fresh produce), it may be desirable to use algorithmic advice to influence the (unknown) risk preferences of decision-makers (DMs) in a risk-averse direction. While there has been considerable analytical work that incorporates risk aversion in inventory models using different objective functions (Ahmed et al., 2007; Prakash Katariya et al., 2014), varying the assumptions regarding the product cost structure (Wang et al., 2009; Wu et al., 2009), exploring different time horizons (Chen et al., 2007) and extending the analysis from single to multiple products (Choi & Ruszczyński, 2011), commensurate behavioral work on risk aversion has been extremely limited. To our knowledge, only a few papers focus explicitly on risk aversion in newsvendor experiments, which represents the inventory management context of the present study. Among these, De Véricourt et al. (2013) observed gender differences in newsvendor ordering behavior, which could be attributed to underlying differences in risk preferences between men and women. Additionally, recent work by Becker-Peth et al. (2018) has demonstrated that risk preferences are significantly correlated with order decisions when the unit of analysis is the individual rather than the group. In this paper, we add to the behavioral newsvendor literature by studying how people with different risk attitudes respond to algorithms of different risk aversion levels.

* Corresponding author.
E-mail addresses: prana.narayanan@novasbe.pt (P. Narayanan), jeeva.somasundaram@ie.edu (J. Somasundaram), matthias.seifert@ie.edu (M. Seifert).

https://doi.org/10.1016/j.ejor.2024.11.013
Received 12 September 2024; Accepted 6 November 2024
Available online 10 November 2024

Additionally, extant studies on decision support in newsvendor experiments do not examine if algorithmic recommendations can influence behavior even *after the algorithm is no longer available* (Feng & Gao, 2020; Zhang & Siemsen, 2019). We believe studying the issue of *algorithmic withdrawal* is important for several reasons. First, in a world where practitioners increasingly resort to licensing third party, off-the-shelf analytics and inventory management software, targeted deployment of algorithmic tools may occur only during limited time periods. For instance, Caro and de Tejada Cuenca (2023) study the adoption of an algorithmic pricing tool that recommends markdown prices during clearance sales campaigns. However, it is crucial for companies to assess whether any behavioral modifications induced by algorithms are likely to persist even after the conclusion of the clearance sales campaign. Moreover, if these algorithm-driven changes do have a lasting impact, then firms may strategically deploy decision support only for a limited duration of time to lower costs. Therefore, it becomes vital for both practitioners and researchers to examine the lasting effects of temporary algorithmic support.

Beyond the field of inventory management, our paper also has broader implications for other resource allocation contexts. For instance, in the domain of finance, subscription-based decision support (also called "robo-advising") is becoming increasingly common not only for professional traders, but also for lay people investing their own private wealth. However, the existing studies on robo-advising have exclusively focused on providing algorithmic recommendations that align with DM's risk preferences (Alsabah et al., 2021; D'Acunto et al., 2019; Niszczota & Kaszás, 2020). In contrast, by exogenously assigning algorithms of different risk-levels, we are able to understand the mechanism by which an individual's risk attitude interacts with the algorithm's risk-level.

Our study also contributes to the decision-making literature on anchoring (Tversky & Kahneman, 1974). In particular, our main theoretical contribution is twofold: First, unlike recent studies that have demonstrated the persistence of anchoring using "uninformative" anchors (Yoon & Fong, 2019), we attempt to show that algorithms that contain informational value regarding a specific level of risk aversion, can leave an imprint on the DM's behavior following temporary exposure. Second, we study the anchoring-and-adjustment mechanism in detail, showing that individuals tend to adjust their order decisions selectively by anchoring more strongly on algorithmic advice that contrasts with their own risk preferences (even after controlling for anchor distance).

Using the anchoring and adjustment heuristic we predict how DMs' order decisions would change in response to algorithmic recommendations. We then test our predictions across two multi-item newsvendor experiments (studies 1 and 2, respectively). We specifically choose the constrained multi-item newsvendor as a cognitively challenging context that is representative of real-world inventory problems compared with the classic single-item newsvendor task (Schweitzer & Cachon, 2000). Additionally, the necessity to allocate capacity (i.e., shelf-space) among multiple alternatives (i.e., products with different demand volatilities) makes the concept of risk salient from the perspective of the DM in the multi-item newsvendor task. Moreover, the experimental setup is similar to that of the portfolio management task, thus making the results generalizable across a wide range of resource allocation problems.

In study 1, we examined the influence of algorithms that have been programmed with varying levels of risk aversion (namely, highly risk-averse (HRA), slightly risk-averse (SRA) and risk-neutral (RN) algorithms) in an inventory task where subjects had to order two products, subject to a shelf-space capacity constraint. We also examined the robustness of our results in the presence of different types of anchors (e.g., human v. algorithm and endogenous vs. exogenous algorithmic assignment). In study 2, we wanted to isolate the effect of risk embedded in the algorithm from the influence of anchor distance (i.e., the difference between the algorithm's risk aversion level and the subject's baseline risk preference). Additionally, we aimed to understand how subject's own risk attitude interacts with that of the algorithm.

Both studies 1 and 2 were comprised of three experimental periods: a baseline period, a treatment period, and a post-treatment period. Each period consisted of several decision rounds. We briefly describe our experimental protocol for study 1. During the baseline period, we studied subjects' order decisions without algorithmic recommendations. Subjects were then randomly assigned either to a control condition, in which they did not encounter any algorithm, or to treatment condition (s), in which subjects observed order recommendations generated by algorithms calibrated with specific levels of risk aversion. Subjects in the treatment condition(s) were aware of the level of risk aversion of the assigned algorithm and were free to decide whether or not to follow algorithmic recommendations. In each decision round, subjects were given extensive feedback on their order decisions, which included insights into whether they had over- or under-ordered relative to the actual demand, their realized profits, and key counterfactuals, such as the profits they would have earned had they precisely adhered to the algorithmic recommendation.[c] Across all conditions, we examined subjects' order decisions during the treatment and the post-treatment period (i.e., after the algorithm was removed).

Our experimental design allowed us to conduct both between-subjects and within-subject comparisons. For instance, we can analyze the order decisions of subjects in the treatment condition(s) in comparison to those in the control group at any given time (between-subjects). Furthermore, we can assess the ordering behavior of subjects in the treatment condition(s) both before and after they encounter the algorithm (within-subject). This distinctive aspect of our research design enhances the reliability of our findings and further distinguishes our study methodologically from existing research on newsvendor decision support (Feng & Gao, 2020; Lee & Siemsen, 2017; Zhang & Siemsen, 2019).

In study 1, we found that highly risk-averse (HRA) algorithmic support led to strong and persistent changes in ordering behavior despite individuals utilizing the HRA algorithmic advice the least. However, the modified (risk-averse) order decisions helped subjects achieve more stable profits. This change in ordering behavior and anchoring effects were similar regardless of whether the HRA algorithm was externally assigned (i.e., exogenous) or chosen by subjects themselves (i.e., endogenous). Additionally, subjects did not exhibit any aversion to highly risk-averse algorithmic recommendations compared to equivalent human advice.

In the follow-up study 2, we found that the risk level programmed in the algorithm is salient to the DM and influences the DM's order decision, beyond the effect of anchor distance. Moreover, we found that subjects adjusted their order decisions by anchoring more on algorithmic advice that contrasted with their inherent risk preferences. For instance, risk-averse subjects tended to rely more on advice that was less risk-averse relative to their baseline decisions, leading them to persistently order in a less risk-averse manner.

Overall, we show that highly risk-averse algorithmic support can be temporarily used to shift multi-item newsvendor order decisions in a risk-averse direction. The anchoring effects are similar regardless of factors such as decision autonomy (endogenous v. exogenous algorithmic assignment) and source of advice (human versus. algorithm). Finally, we show that subjects selectively adjust their order decisions by anchoring more on algorithms that contrast their inherent risk preferences in the newsvendor task.

Our study is rooted in decision-making and behavioral operations literature. Specifically, we draw on past research related to the anchoring and adjustment heuristic to examine behavioral responses in a classic resource allocation problem faced by operations managers involving inventory ordering decisions.

[c] The counterfactual was only displayed to subjects in the algorithmic advice conditions during the treatment period.

1.1. Anchoring and adjustment

Anchoring refers to the tendency of DMs to be influenced in their judgments by initially presented values (Chapman & Johnson, 2012; Tversky & Kahneman, 1974). Anchoring has attracted extensive scholarly attention partly due to its broad applicability across diverse domains (Beggs & Graddy, 2009; Englich et al., 2005; Englich & Mussweiler, 2001; Fujiwara et al., 2013; Loschelder et al., 2016; McAlvanah & Moul, 2013; Wegener et al., 2010). In behavioral operations, Schweitzer and Cachon (2000) used anchoring on mean demand as a possible explanation for the Pull-to-Center bias in the single item newsvendor problem, which is the tendency of decision-makers (DMs) to place orders that lie between mean demand and optimal order quantity. This account has recently been corroborated from the perspective of prospect theory (Uppari & Hasija, 2019).

A common finding from extant studies that use the anchoring paradigm is that advice affects subsequent judgement if it is presented before the DM has made an independent estimate (Koehler & Beauregard, 2006; Rader et al., 2015; Sniezek & Buckley, 1995). However, extant models of anchoring do not offer any predictions about the persistence of anchoring effects even after the anchor is removed (Lieder et al., 2012; Turner & Schley, 2016).

The anchoring paradigm we propose in this paper has two key features: First, to predict treatment period effects, we assume that DMs anchor their order decisions based on the algorithm's recommendation and insufficiently adjust towards an optimal level that is based on their own baseline order decisions (i.e., anchoring-and-adjustment heuristic). Second, we allow for changes in ordering behavior to persist in the post-treatment period, which may be due to automaticity of order decisions (e.g., see habit formation models - Becker & Murphy, 1988). We first describe anchoring below in the context of an inventory management setting (also referred to as multi-item newsvendor problem in the operations literature).

1.2. Inventory management problem

1.2.1. Setup

A DM in our resource allocation context sells two kinds of perishable, non-substitutable products: A and B. Both products have identical cost structures (i.e., unit selling price of p and unit cost price of c where $p > c$) and are *high-profit* margin products, meaning the *critical fractile* for either product is > 0.5 (Schweitzer & Cachon, 2000). The demand for products A and B is uncertain and varies across time periods. Let d_{At} and d_{Bt} denote the actual demand for products A and B in a particular time period t. The actual demand values are drawn from two different, discrete uniform distributions, which are known to the DM at the start of time period t. Additionally, the DM is aware that demand values of products A and B are independent (both within time period t and across time periods). Let μ_{At} and μ_{Bt} denote the mean demand for products A and B and σ_{At} and σ_{Bt} represent the standard deviation of demand for products A and B for time period t. In our task setup, product B's demand is more volatile than that of product A. More specifically, the coefficient of variation of demand for product B was about 2 times the coefficient of variation of demand for product A - i.e., $\frac{\sigma_{Bt}}{\mu_{Bt}} = 2 * \frac{\sigma_{At}}{\mu_{At}} \ \forall \ t$.

The DM has to decide the order quantities of products A and B (denoted by q_{At} and q_{Bt}, respectively) in each time period t, while adhering to a resource constraint on the store shelf-space capacity (i.e., total units of product A and B ordered must be less than or equal to 100). Importantly, actual demand values are realized only after the order decision is made. As in the case of a single-item newsvendor, the DM could either under- or over-order either product. There are no salvage costs or stockout costs included in our multi-item newsvendor setup.

Random profits for each product (denoted by π_{At}, π_{Bt}) in time period t are calculated using Eqs. (1) and (2) below. The total random profit (π_t) is the sum of the individual profits of product A and B.

$$\pi_{At} = p\min(q_{At}, d_{At}) - cq_{At} \tag{1}$$

$$\pi_{Bt} = p\min(q_{Bt}, d_{Bt}) - cq_{Bt} \tag{2}$$

$$\pi_t = \pi_{At} + \pi_{Bt}. \tag{3}$$

Algorithm design. We design the algorithm such that it tries to maximize an objective function F, which can be expressed as follows:[d]

$$F = E(\pi_t) - \lambda * V(\pi_t) \tag{4}$$

In Eq. (4), $E(\pi_t)$ is the expected profits, $V(\pi_t)$ represents the variance in profits and λ captures the degree of risk aversion of the algorithm. The intuition behind Eq. (4) is that an increase in risk aversion is modeled by an increasing distaste for payoff variance. Therefore, *higher values of risk aversion λ should lead to less ordering of the riskier product B.* For the study parameters we used, we simulated this result, and it is shown in Online Appendix 1.

The optimal product B order quantities suggested by the RN, SRA and HRA algorithms are denoted by q^*_{RN}, q^*_{HRA} and q^*_{SRA} respectively. For study 1, we chose λ values of 0.5 and 2 to model the SRA and HRA algorithms, respectively. Note that $\lambda = 0$ for a RN algorithm.[e] We confirmed that $q^*_{RN} > q^*_{SRA} > q^*_{HRA}$ in study 1 for all decision rounds.

1.2.2. Predictions

The objective is to examine how DMs modify their order decisions in response to the algorithmic recommendations and derive predictions for our experiment., We focus on the three main periods of the experiment: (i) Baseline period $(t = 0)$ where the DM makes order decisions without any algorithmic aid, (ii) Treatment period $(t = 1)$, in which DMs (randomly) assigned to treatment conditions make order decisions after observing algorithmic recommendations of a specific risk aversion level (HRA, SRA or RN) and (iii) Post-treatment period $(t = 2)$, where the DM continues to make the order decisions, but without any algorithmic support. Note that subjects (randomly) assigned to a *control* condition do not encounter any algorithmic recommendation in any of the three periods. Each experimental period consists of several decision rounds. We allow q_B (henceforth q) in our set-up to be influenced by external anchors (e.g., algorithmic recommendations) and vary over time t. This allows us to derive a unique set of predictions about the order decision in each experimental time period. We assume that q_t evolves as follows:

1.2.3. Baseline period

As subjects are randomly assigned, we do not expect to see any difference in (average) q_0 across conditions at the end of the baseline period i.e., $q_0^{HRA} = q_0^{SRA} = q_0^{RN} = q_0^{Control} = q_0$.

1.2.4. Treatment period

During the treatment period, the DM's order decision could be influenced by an external anchor in the form of an algorithmic recommendation. We assume that in the presence of an algorithm, the DMś final order is a weighted average of their baseline order q_0 and the algorithmic recommendation. In the multi-item newsvendor task, we

[d] The objective function of the form $F = E - \lambda V$ is identical to the mean-variance framework typically used in the classical portfolio management task and can be derived by assuming a DM tries to maximize an exponential utility function with a particular level of risk aversion λ and a normal distribution of profits (Sargent, 2009)). While we note that the newsvendor profits are not normally distributed, this formulation of F suffices to generate intuitively explainable risk-averse algorithmic recommendations (For an illustration, refer to Online Appendix 1).

[e] q^*_{RN} can also be obtained in every time period from the heuristic proposed by Erlebacher (2000), where $q^*_{RN} = \mu_{Bt} + \frac{\sigma_{Bt}}{\sigma_{At}+\sigma_{Bt}}(100 - \mu_{At} - \mu_{Bt}).\langle\underline{END}\rangle$We lack an empirical benchmark in the multi-item newsvendor task to determine whether the baseline order (q_0) would be closer to q^*_{RN}, q^*_{SRA} or q^*_{HRA}.

expect that the HRA algorithmic suggestion (q^*_{HRA}) to be the anchor that is possibly the furthest away from subjects' baseline order quantity (compared to SRA and RN algorithmic suggestions) i.e., $q_0 \gg q^*_{HRA}$.[f] According to advice-taking literature, decision makers (DMs) tend to deviate the most from advice that is furthest from their own estimates, a phenomenon more broadly known as egocentric discounting (Rader et al., 2017; Yaniv & Kleinberger, 2000). However, due to the extremity of the HRA algorithmic anchor, we anticipate that the adjustment will still be insufficient—meaning the order decisions will not fully return to baseline levels. Therefore, during the treatment period, we hypothesize that subjects exposed to the HRA algorithm will order relatively fewer units of the riskier Product B compared to those who observe SRA or RN algorithmic recommendations, which are expected to be closer to subjects' baseline order decisions.

H1. DMs who observe a highly risk-averse algorithm order less of the riskier product (B) compared to those who do not observe any algorithm or those who observe slightly risk-averse or risk-neutral algorithm.

1.2.5. Post-treatment period

Note that in the post-treatment period no anchor (algorithmic recommendation) is provided to DMs. However, we suggest that order decisions during the treatment period may still impact order decisions in the post-treatment period due to the persistence of previously encountered anchoring effects. While anchoring has been extensively studied across multiple domains, there is relatively little empirical work on the temporal persistence of anchoring, especially when the anchor is externally provided. In a notable exception, Yoon and Fong (2019) find that uninformative external anchors (i.e., random values uncorrelated with the true estimate) still have a lasting effect on valuation judgments, such as DMs' willingness-to-pay (WTP). While recent research suggests that relevant anchors (such as the algorithmic recommendation we use) are expected to have larger effects compared to irrelevant anchors, the results are still mixed, especially regarding extreme anchors (Glöckner & Englich, 2015). The persistent effect of relevant anchors would be broadly in line with theories on habit formation (Becker & Murphy, 1988). According to these theories, the key feature driving behavioral stickiness is referred to as *adjacent complementarity,* i.e., the extent to which a particular behavior has been repeated in the past. In our case, this implies that subjects who ordered less of the riskier product B during the treatment period are likely to continue to do so during the post-treatment period. Thus, based on findings from prior literature, we expect anchoring effects to persist and hypothesize the following:

H2. DMs who observe a highly risk-averse algorithm continue to order less of the riskier product (B) even after the algorithm is removed.

We test our predictions in study 1. Notably, H1 and H2 are only based on anchor distance (i.e., how far the algorithmic recommendation is from the subject's baseline order). Later, in study 2, we will examine the effect of risk embedded in the algorithmic recommendation while experimentally controlling for the anchor distance. All data and code from studies 1 and 2 are publicly available on ResearchBox (Link).

2. Study 1

2.1. Methods

2.1.1. Subjects

Subjects were Master's degree students from a large private university. Study 1 was conducted online, and subjects earned class credits in return for their participation. In addition to earning extra course credits, we selected 10 subjects at random and awarded them Amazon vouchers matching the profits of one randomly chosen round.[g] The average earnings were $22 per winning subject. Our final sample size was 184 (female = 40 %) with an average age of 27 years.

2.1.2. Design and procedure

Once subjects consented to take part in the study and were briefed on the incentives, they were asked to assume the role of a manager of a retail store which sold products A and B. Subjects were informed about the multi-item news vendor set-up described in Section 2.2 and were asked to order a total of 100 units, since the shelf-space capacity constraint was binding for any level of risk aversion. The complete set of task instructions used in study 1 is provided in Online Appendix 2. The demand distributions for products A and B are listed in Table 1 below.

In Table 1, z_t is a random integer between 0 and 12, independently generated for every round t (and undisclosed to subjects). We incorporated random parameter z_t to vary the demand distribution in each round, so that subjects encountered a (slightly) different algorithmic recommendation every time. However, the demand distribution (i.e., z_t value) and actual demand values realized were identical for all subjects in any given round of the study. Based on the range of possible demand values, subjects were also informed that the demand volatility of product A (σ =10.1) was lower than the demand volatility of product B (σ =36.1). Also, as the coefficient of variation of product B's demand was greater than that of product A, subjects were told that product B was *riskier* than product A (i.e., ordering more of product B would result in more volatile profits across rounds).

Consistent with our model setup, the study comprised three periods: a baseline period (8 rounds), a treatment period (20 rounds) and a post-treatment period (12 rounds). In every round, subjects observed the demand distribution of products A and B and decided how many units to order of each product, subject to the shelf-space capacity constraint. At the end of the baseline period, subjects were randomly assigned to one of four experimental conditions: a control condition, in which they did not encounter any algorithm, or to one of three treatment conditions in which they observed recommendations generated by a highly risk-averse (HRA), slightly risk-averse (SRA) or a risk-neutral (RN) algorithm. In each of the treatment conditions, subjects were informed about the level of risk associated with the algorithmic recommendation and were free to decide whether or not to follow the algorithmic recommendation. Fig. 1 below summarizes the experimental design for study 1.

At the end of each round, all subjects received two types of feedback: one in the form of a summary table and the other in the form of a graph. In tabular form, subjects were reminded of their chosen order quantities, the algorithmic recommended order quantities, the actual demand values, and whether or not they under- or over-ordered each product.[h] In the graph, subjects viewed their actual profits along with two counterfactuals: the profits they would have received had they followed the algorithmic recommendation, and the maximum possible profits. Maximum possible profits were computed assuming that the order quantities matched the realized demand values. An example of the feedback screen displayed to study participants is shown in Fig. 2 below.

Table 1
Demand distribution of products A and B.

Product type	Range of possible demand values
A	$[(15 + z_t), (50 + z_t)]$
B	$[(0 + z_t), (125 + z_t)]$

[f] This variation of the random incentive system is almost the exclusively used incentive system in similar individual choice experiments involving risk preferences (van de Kuilen & Wakker, 2011)

[g] The algorithmic recommended order quantity was only displayed to subjects in the treatment conditions condition during the treatment period.

[h] In Fig. 3, we subtracted the random number z_t from every subject's order.

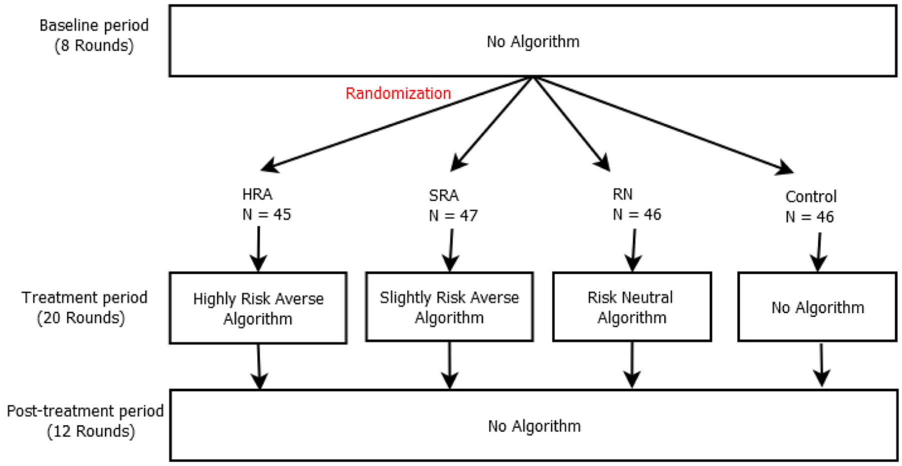

Fig. 1. Study 1 design.

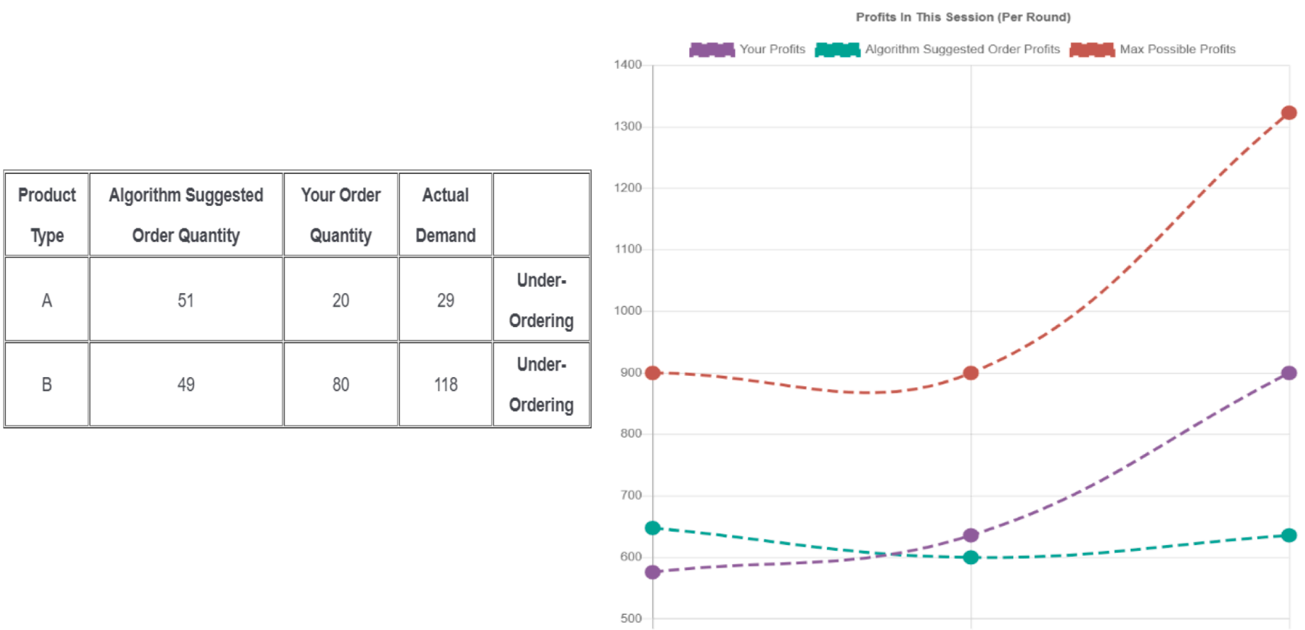

Product Type	Algorithm Suggested Order Quantity	Your Order Quantity	Actual Demand	
A	51	20	29	Under-Ordering
B	49	80	118	Under-Ordering

Fig. 2. Feedback screen displayed to a hypothetical participant in study 1.

At the end of the treatment period, we removed the algorithmic aid for all subjects, who then proceeded to complete the final twelve rounds of the study in the post-treatment period. At the end of the task, subjects reported their demographic information (age and gender) and answered a few survey questions related to their decision-making strategy.

3. Results

3.1. Ordering behavior

We operationalized ordering behavior in study 1 as the number of units of (riskier) product B ordered per round. We analyzed ordering behavior in each of the three experimental periods separately. In Fig. 3, we observe no discernible differences in ordering behavior among subjects across different conditions during the baseline period due to random assignment (also see Table 2, columns 1–2) .[i] In addition, we see that subjects in all three treatment conditions ordered fewer units of product B than those in the control condition during the treatment period (also see Table 2, columns 3–4). For example, subjects in HRA, SRA and RN conditions ordered, on average, 7.4, 4.9 and 3.8 fewer units of product B, respectively, compared to subjects in the control condition during the treatment period. In Fig. 3, we again observe systematic differences in ordering behavior among subjects based on their assigned experimental condition during the post-treatment period. Despite removing the algorithm in the post-treatment period, subjects in the HRA condition continued to order fewer units of product B compared to those in the control condition (also see Table 2, columns 5–6). While the effects marginally diminished from treatment to post-treatment period, they were still significant ($p < 0.01$). For example, subjects in HRA condition ordered, on average, 4.6 fewer units of product B compared to those in the control condition during the post-treatment period (between-subjects). Additionally, a paired *t*-test (within-subjects) revealed that subjects in the HRA condition ordered, on average, about 5 fewer units of product B in the post-treatment period compared to the baseline period ($t(44) = 6.21$, $p < 0.001$). Thus, study 1 provides support for H1

[i] Additional robustness check to examine post-treatment stickiness is presented in Online Appendix 3.

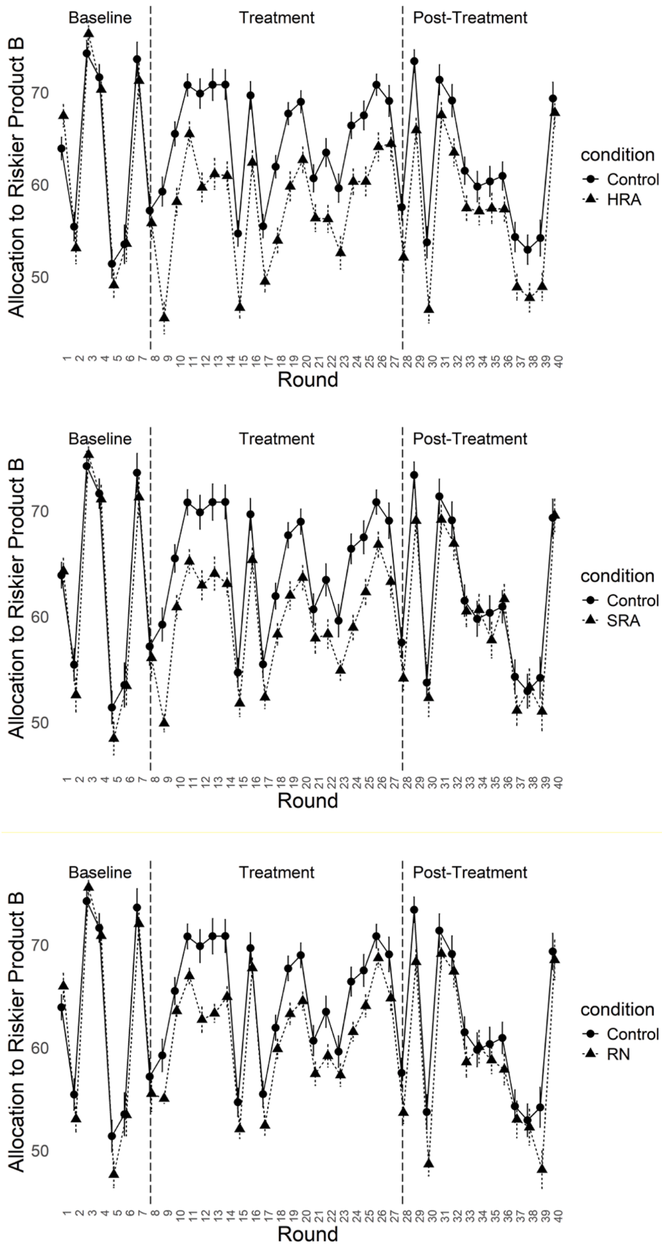

Fig. 3. Allocation to riskier product B
Note: Error bars signify ±1 standard error.

and H2.[j]

3.1.1. Advice utilization

Although subjects in the HRA condition changed their order decisions the most, they also showed the largest deviation from algorithmic recommendation. Subjects in the HRA condition ordered, on average, 11 units more of product B than what was suggested by the HRA algorithm. Subjects in the SRA condition also deviated from the algorithmic recommendation, but the magnitude of deviation was relatively smaller (See Online Appendix 3). Consistent with prior literature, we also computed a weight on advice (WOA) score to capture the extent of advice utilization (Logg et al., 2019; Soll & Larrick, 2009). Weight on advice (WOA) is calculated as follows:

$$WOA = \frac{final\ estimate - initial\ estimate}{advice - initial\ estimate} \tag{5}$$

Higher values of WOA indicate greater advice utilization. For instance, WOA = 0 implies that there is no revision of the initial estimate based on the advice (i.e., 0 % advice utilization). Conversely, WOA = 1 indicates that the final estimate matches the advice provided (i.e., 100 % advice utilization). WOA values <0 or greater than 1 are possible, but unlikely because final estimate mostly tends to lie between the initial estimate and advice. In line with previous research, we re-coded any negative WOA values as zeros, and any values above 1 as 1. In our experiment, we approximated the initial estimate as the subject's average baseline order (since we did not explicitly ask for initial estimates) and a WOA score was calculated for each subject in the treated condition in every round of the treatment period as follows:

$$WOA = \frac{q - q_{baseline}}{q_{advice} - q_{baseline}} \tag{6}$$

where q is the subject's order of product B, q_{advice} is the corresponding algorithmic recommendation and $q_{baseline}$ is the initial estimate (made without the algorithmic recommendation). We found that subjects who encountered the RN algorithm utilized the advice 24 % more than those who observed the HRA algorithm ($\overline{WOA}_{HRA} = 0.34$; $\overline{WOA}_{RN} = 0.58$). The results are shown in Fig. 4 below.

Additionally, following recommendations from recent work (e.g., Himmelstein & Budescu, 2023; Rebholz et al., 2024) we also computed two variants of the classic WOA measure to understand how the algorithmic advice shifts the distribution of the initial advice. These procedures yield similar results and are reported in Online Appendix 3.

3.1.2. Average profits

The average profits earned by subjects in all conditions during each of the three experimental periods are reported in Online Appendix 3. Subjects in RN and SRA conditions earned average higher profits per round ($24 and $13 respectively) than those in the HRA condition during the treatment period ($t(73.98) = 5.39$, $p < 0.001$ and $t(87.64) = 2.92$, $p = 0.004$, respectively). However, we did not find any significant difference in average profits earned by subjects across different conditions during the post-treatment period.

3.1.3. Variability in profits

We measured variability in profits as the average standard deviation of profits, computed for each subject across treatment and post-treatment periods, separately. From Fig. 5 below, we see that subjects in the HRA Algorithm condition experienced lower variability in profits compared to those in the RN Algorithm and control conditions in the treatment period. Additionally, subjects in the HRA Algorithm condition continued to experience lower variability in profits compared to those in the control condition in the post-treatment period (See, Fig. 5 below).

Overall, we find that subjects changed their order decisions in the risk averse direction strongly and persistently due to the influence of HRA algorithm. This also resulted in more stable profits compared to subjects in other conditions. We explore the underlying reasons for the change in order decisions for subjects in HRA condition in sub-Section 2.3.[k]

3.1.4. Persistence of treatment effects

We examine the reasons for change in order decisions in the post-treatment period compared to the baseline period. Before we test our anchoring paradigm, we consider various alternate explanations,

[j] We do not examine behavioral biases since order decisions for each product are not independent.

[k] 54% of subjects said they would choose the algorithm over human advice in a similar task in the future whereas only 24% of subjects said they would choose the human advisor over the algorithm. The remaining 22% were indifferent between the two options (human v. algorithmic advice).

Table 2
OLS regression on ordering behavior with round-fixed effects included.

Standard errors were clustered at the subject-level and are reported in parentheses.

	Dependent variable:					
	Allocation to Riskier Product B					
	Baseline		Treatment		Post-Treatment	
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	65.992***	65.913***	64.559***	39.937***	71.327***	42.637***
	(1.046)	(3.364)	(0.988)	(5.493)	(1.050)	(5.981)
Highly Risk-averse	−0.486	−0.578	−7.415***	−7.455***	−4.582***	−4.647***
	(1.259)	(1.267)	(1.074)	(0.934)	(1.201)	(1.115)
Risk-neutral	−0.852	−0.960	−3.870***	−3.808***	−2.513**	−2.445**
	(1.202)	(1.209)	(0.833)	(0.729)	(1.186)	(1.082)
Slightly Risk-averse	−1.042	−1.081	−5.216***	−4.932***	−1.496	−1.169
	(1.246)	(1.231)	(0.933)	(0.796)	(1.367)	(1.281)
Age		−0.006		0.013		0.065
		(0.120)		(0.086)		(0.110)
Male		0.484		1.105*		1.346
		(0.892)		(0.618)		(0.878)
Baseline Allocation				0.346***		0.383***
				(0.068)		(0.072)
Observations	1472	1472	3680	3680	2208	2208
Adjusted R^2	0.139	0.138	0.146	0.190	0.066	0.108
Residual Std. Error	10.461	10.465	8.671	8.442	10.404	10.164

*Note:**$p < 0.1$; **$p < 0.05$; ***$p < 0.01$, Reference Group: Control.

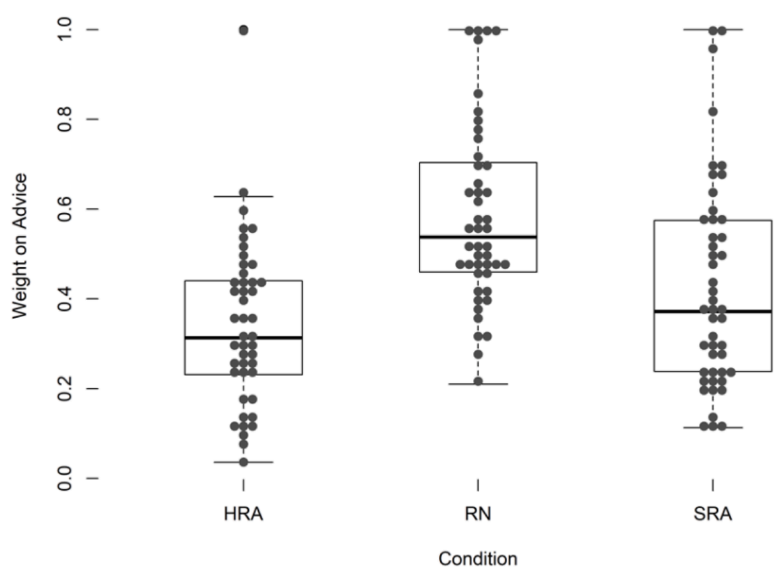

Fig. 4. WOA score distribution
Note that each black dot represents the average WOA computed for a subject during treatment period.

including the change in underlying risk preferences, loss aversion and implicit anchoring on mean demand heuristic. However, these explanations fail to account for the observed effects in the post-treatment period (See Online Appendix 4). Nonetheless, if anchoring effects persist, we anticipate that individuals who are more influenced by the algorithmic anchor during the treatment period will subsequently order less in the post-treatment period. Such a finding would be broadly in line with theories on habit formation (Becker & Murphy, 1988).

In order to test this proposition, we measured the correlation between the "change in the average order of product B" during the post-treatment period (vis-à-vis baseline) $(\overline{q}_2 - \overline{q}_0)$" and the "change in the average order of product B during the treatment period (vis-à-vis baseline) $(\overline{q}_1 - \overline{q}_0)$" for subjects in the HRA condition in study 1. From Fig. 6 below, we see that the order decisions in the treatment and post-treatment periods correlate positively subjects in the HRA condition ($r = 0.66,\ p < 0.01$) thus supporting the persistence of the anchoring hypothesis. A more sophisticated instrumental regression (IV) analysis is reported in Online Appendix 5.

4. Robustness of the results to different anchors

Since previous research suggests there could be differences in algorithmic utilization depending on whether the algorithmic is endogenously or exogenously assigned (Dietvorst et al., 2018; Rader et al., 2017a) and whether the recommendations originate from an algorithm or a human (Dietvorst et al., 2015), we briefly examine whether the findings of our study may have been influenced by decision autonomy (i.e., choosing an algorithm versus random assignment) and/or source of advice (human versus algorithm).

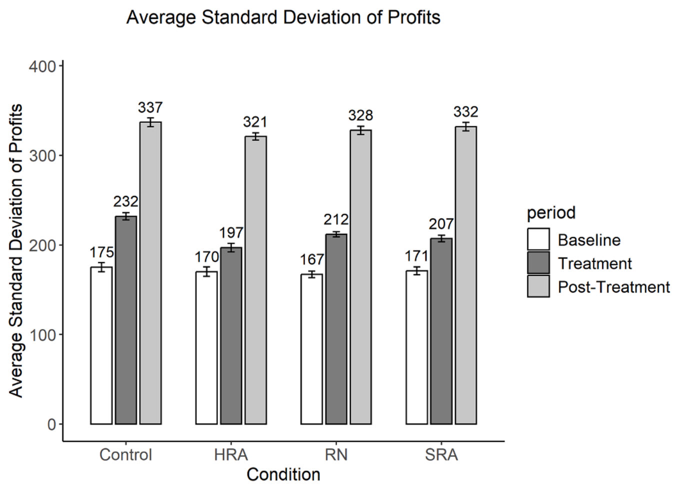

Fig. 5. Study 1: Average standard deviation in profits
Note: Error bars signify ±1 standard error.

a human or an algorithm. Thus, we sought to understand whether the source of advice influences the strength and persistence of anchoring effects in a pre-registered lab experiment (Link). We compared the HRA algorithm recommendation with equivalent human advice (from our study 1) that exhibited a similar level of risk aversion. Details regarding the methods and results of this supplementary study are reported in Online Appendix 7. While subjects in HRA algorithm condition ordered slightly fewer units of riskier product B than those in the HRA human condition during the treatment period, the difference between the two conditions was not significant during the post-treatment period. However, in a post-experimental survey, we found evidence for algorithmic appreciation.[1]

5. Study 2

In study 1, we demonstrated the lasting influence of highly risk-averse algorithmic anchors in the multi-item newsvendor context. However, the treatment and post-treatment effects in study 1 might have been driven solely by the distance of the algorithmic anchor from the DM's baseline orders. To isolate the effect of risk attitude embedded in the algorithmic advice from the impact of anchor distance, in study 2, we not only explicitly inform subjects about their own risk attitude based on their ordering behavior during the baseline period (compared to the median risk attitude in the subject pool) but also provide them with a customized algorithm—manipulated to be either risk-averse or risk-seeking — that was set at a similar distance from each DM's inherent risk attitude. This approach allowed us to control for anchor distance and isolate the effect of risk attitude on the shift in order decision behaviors. Moreover, this design helps to address how individuals with different risk attitudes respond to algorithmic recommendations that vary in their risk levels, thus offering a more detailed and nuanced account of the anchoring-and-adjustment mechanism.

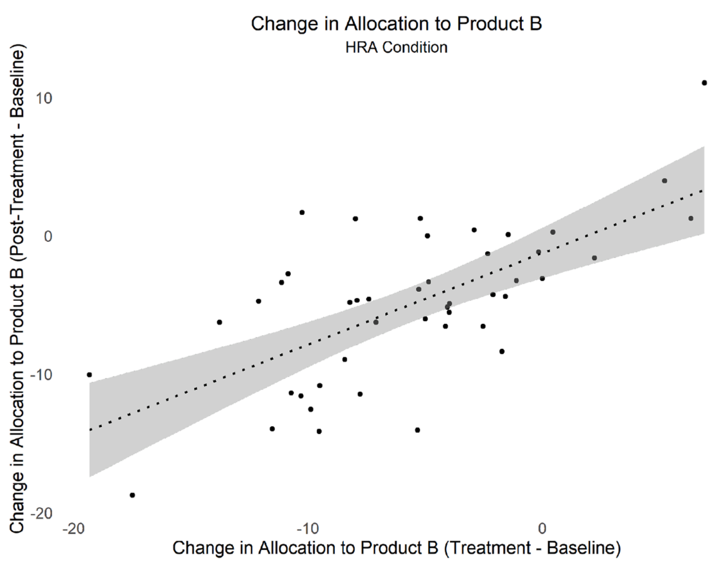

Fig. 6. Change in allocation to riskier product B: HRA condition
Note: Each dot represents the change in the average order of Product B for a subject in the HRA condition. The gray shaded region represents 95 % confidence interval estimate.

4.1. Decision autonomy (Exogeneous vs endogenous)

Understanding the role of *decision autonomy* is relevant for practical purposes, especially for organizations where managers have the possibility to select an algorithm with a particular risk preference. In fact, the practice of matching DMs with an algorithm based on a specific risk aversion is common in similar resource allocation contexts (D'Acunto et al., 2019). Therefore, we sought to explore whether anchoring effects are stronger and more persistent when DMs have the *autonomy* to select an algorithm with a predetermined level of risk aversion, as opposed to having the algorithm assigned to them at random (i.e., as in study 1). We examined this in a pre-registered lab experiment (Link) and the methods and results are reported in Online Appendix 6. Importantly, we found no differences in ordering behavior or algorithmic uptake based on whether the algorithm was chosen by the subject or randomly assigned.

4.2. Source of advice (Human vs algorithm)

Besides decision autonomy, an important practical consideration for any type of decision support is whether the recommendations come from

[1] The median value of 68.5 was computed from the original newsvendor study 1 and it worked out reasonably well, in that we had a roughly 50-50 split of Risk-Averse and Risk-Seeking subjects in study 2.

5.1. Methods

5.1.1. Subjects

As in study 1, subjects were Master's degree students. Study 2 was conducted online, and subjects earned class credits in return for their participation. In addition to earning extra course credits, we selected 10 % of subjects at random and awarded them Amazon vouchers based on profits earned in a randomly chosen round. The average earnings were $6 per winning subject. Our final sample size was 114 (female = 59 %) with an average age of 28.3 years. The study was pre-registered on AsPredicted (Link).

5.1.2. Design and procedure

Study 2 was designed to closely resemble study 1 in all respects, except for two key distinctions. First, at the end of the baseline period, subjects in study 2 were explicitly informed whether they were *Risk-Averse* or *Risk-Seeking* (relative to the median risk attitude in the subject pool) that was inferred from their average order of riskier product B during the baseline period. Specifically, participants who ordered more or less than the median value (68.5) of Product B during the baseline period were informed to be either *Risk-Seeking* or *Risk-Averse* compared to the median subject.[m] Second, subjects were randomly assigned to one of two possible algorithmic advice treatments: in the less risk-averse condition (*LRA*), participants observed algorithmic recommendations that were programmed to be systematically less risk averse than themselves (or their baseline orders), whereas in the more risk-averse condition (*MRA*), they were assigned to an algorithm generating recommendations that were systematically more risk averse than themselves (or their baseline orders). Importantly, *LRA* and *MRA* recommendations were tailored for each participant such that the risk-levels of the algorithmic recommendations were adjusted from the baseline risk-level by subtracting or adding the same constant (0.5). For instance, subjects whose baseline risk-level was $\lambda = 0.2$ would either face a *LRA* or *MRA* algorithm — whose λ was either - 0.3 or 0.7, respectively — with equal chance. This manipulation allows us to fix the effect of anchor distance and isolate the impact of algorithmic risk on order decisions. We again used the objective function $E - \lambda V$ to map any risk-aversion level (λ) to a particular allocation to product B (q_B) and generate the algorithmic recommendations.

To summarize, in study 2, we inform subjects of their relative risk attitude (i.e., *Risk-Averse* or *Risk-Seeking*) and randomize them into two different treatments that only differed in the level of risk of the algorithmic recommendations (i.e., More or Less Risk-Averse relative to the subject's baseline orders). The main purpose of this experimental design is to understand how people with different risk-attitudes respond to algorithms with different risk-levels (while anchor distance is fixed). As the demand parameters were identical across all rounds, we use the control condition from study 1 as a reference group to examine post-treatment stickiness.[n]

5.2. Results

5.2.1. Ordering behavior

We again operationalized ordering behavior in study 2 as the number of units of (riskier) product B ordered per round. We analyzed changes in ordering behavior from the baseline period and ran a regression analysis in Table 3 using the control group from study 1 as the reference group. We find evidence for selective changes in ordering behavior, which occurs when individuals of a particular risk attitude face algorithms whose risk preferences are in the opposite direction. For instance, *Risk-Seeking* subjects who encountered the *MRA algorithm* anchored on the advice subsequently became more risk-averse by ordering 6 % and 5 % fewer units of product B (on average, relative to the control group) during the treatment and post-treatment periods, respectively. Similarly, *Risk-Averse* subjects who faced the *LRA algorithm* became more risk-seeking by ordering 11 % and 7 % more units of product B (on average, relative to the control group), during the treatment and post-treatment periods, respectively. On the other hand, *Risk-Seeking* subjects who faced LRA algorithm or Risk-*Averse* subjects who faced the MRA algorithm did not change their order decisions.[o]

5.2.2. Advice utilization

We used the weight on advice (WOA) measure reported earlier in study 1 as a measure of advice utilization. The regression results are reported in Online Appendix 8. We find that LRA algorithmic advice is utilized 16 % more (on average) than the MRA advice (See Fig. 7 below). Moreover, our findings on advice utilization align with the results on order decisions. Specifically, *Risk-Averse* subjects tend to rely more on the LRA advice, while *Risk-Seeking* subjects anchor more on the MRA advice, despite both recommendations being at a similar distance from the subject's baseline risk level.

Overall, in study 2, where subjects are informed about their own risk attitude and we experimentally control for the distance of the anchor,

Table 3
OLS regression on change in order decisions (from baseline) with round-fixed effects included.

Standard errors were clustered at the subject-level and are reported in parentheses

	Dependent variable:			
	% Change in Allocation to Riskier Product B from Baseline			
	Treatment		Post-Treatment	
	(1)	(2)	(3)	(4)
Constant	0.010	−0.010	0.030	−0.036
	(0.018)	(0.070)	(0.020)	(0.079)
LRA Condition - Risk Averse	0.114***	0.113***	0.077*	0.077*
	(0.029)	(0.029)	(0.045)	(0.045)
LRA Condition - Risk Seeking	−0.011	−0.008	−0.024	−0.022
	(0.021)	(0.021)	(0.025)	(0.025)
MRA Condition - Risk Averse	0.056*	0.054*	0.102**	0.096**
	(0.033)	(0.032)	(0.043)	(0.042)
MRA Condition - Risk Seeking	−0.064***	−0.063***	−0.051**	−0.051**
	(0.019)	(0.019)	(0.025)	(0.025)
Age		0.0003		0.002
		(0.002)		(0.003)
Male		0.029*		0.025
		(0.015)		(0.022)
Observations	3120	3120	1872	1872
Adjusted R^2	0.128	0.134	0.069	0.073
Residual Std. Error	0.173 (df = 3096)	0.172 (df = 3094)	0.208 (df = 1856)	0.208 (df = 1854)

*Note:**$p < 0.1$; **$p < 0.05$; ***$p < 0.01$, Reference Group is Control.

[m] The lack of an explicit control condition (due to sample size constraints) is the only notable deviation from our pre-registration report for study 2 (Lakens, 2024).

[n] Surprisingly, during the post-treatment period, *Risk-Averse* subjects shifted towards risk-seeking behavior, even when presented with a *MRA algorithm*, by deviating away from its recommendations. We believe this behavioral pattern is likely to result from a framing effect, in that *Risk-Averse* individuals may prefer not to be perceived as overly cautious and thus hedge in a risk-seeking direction, regardless of the advice given.

[o] Surprisingly, during the post-treatment period, *Risk-Averse* subjects shifted towards risk-seeking behavior, even when presented with a *MRA algorithm*, by deviating away from its recommendations. We believe this behavioral pattern is likely to result from a framing effect, in that *Risk-Averse* individuals may prefer not to be perceived as overly cautious and thus hedge in a risk-seeking direction, regardless of the advice given.

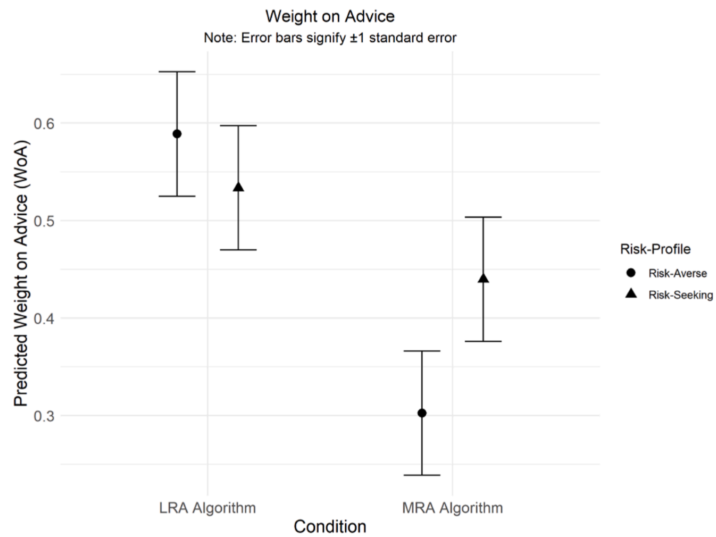

Fig. 7. Weight on advice (WOA) Score by condition and risk-profile.

we observe a clear pattern of selective adjustment in ordering behavior. Specifically, *Risk-Averse* individuals tend to anchor more on the LRA algorithmic advice, leading them to order more units of the riskier Product B. This effect is largely symmetrical, as *Risk-Seeking* subjects exhibit a similar pattern by ordering fewer units of Product B when exposed to the MRA algorithmic recommendation. Crucially, study 2 documents the influence of risk attitude embedded in the algorithm beyond the effect of anchor distance.

6. Discussion and conclusion

We study individual decision-making in resource allocation tasks when temporary algorithmic recommendations of varying risk levels are available. Using the literature on anchoring, we predicted that DMs' order decisions change when they are temporarily exposed to an algorithmic aid. We tested the empirical validity of our claims across a series of experiments that are based on a classic inventory management problem. In study 1, we showed that DMs in a multi-item newsvendor task ordered less subsequent to observing recommendations from a highly risk-averse algorithm. Importantly, DMs continued to order less even after the algorithmic recommendation was removed. The modified order decisions helped reduce the variability in profits. Furthermore, we showed that DMs did not exhibit any algorithm aversion and advice utilization was the same regardless of whether the algorithm was assigned externally or chosen by the subjects themselves. Finally, in a follow-up experiment (study 2), we demonstrated that risk level embedded in the algorithm is highly salient to DMs and subjects tended to rely more on algorithmic advice that contrasted with their baseline risk preferences.

In study 1, a key observation is that subjects' order decisions were most influenced by highly risk-averse (HRA) algorithmic recommendations, even though they placed less weight on the HRA algorithmic advice (as indicated by the low WOA score). Thus, the least used advice resulted in the most significant changes in behavior. How do we explain this? First, from an experimental design perspective, the low levels of advice utilization for HRA algorithm can be attributed to the fact that subjects were fully aware of the algorithm's high level of risk aversion and had access to detailed feedback on profits, which were low (on average) for highly risk-averse order recommendations. Second, based on the literature on advice-taking, we know that DMs tend to deviate most from advice that are farthest away from their own estimates, a phenomenon more broadly referred to as egocentric discounting (Rader et al., 2017; Yaniv & Kleinberger, 2000). Thus, given that HRA algorithmic recommendation represents the most extreme type of advice (or anchor), it makes sense that subjects would deviate substantially from it. However, due to the extremity of the HRA anchor, the deviation was insufficient—meaning the order decisions did not fully return to baseline levels. This explains why we observe significant changes in ordering behavior for subjects who encountered the HRA algorithm, despite the low levels of advice adherence. Additionally, we explain the persistence of change in ordering behavior for DMs who observe HRA algorithms as a consequence of anchoring and repetition (Becker & Murphy, 1988). In our case, this implies that subjects who ordered less of the riskier product in the treatment period continue to do so during the post-treatment period. Finally, we do not observe strong changes in ordering behavior for subjects who observed the SRA and RN algorithms, because the recommendations were already quite close to subjects' baseline orders. Thus, the advice had little impact on behavior (Soll et al., 2022).

In study 2, the key takeaway is that anchoring effects are not merely driven by the extremity of the anchor. When the anchor distance was held constant and participants were exposed to algorithms with varying levels of risk aversion, we observed evidence of selective adjustment in order decisions, with individuals anchoring more strongly on algorithmic advice that contrasted with their inherent risk preferences. Risk-averse participants exposed to the LRA algorithm became more risk-seeking, while risk-seeking participants who encountered the MRA algorithm became more risk-averse. On the other hand, participants who faced an algorithm which was similar to their baseline risk-preferences did not change their order decisions. We interpret these findings as indicative of a subject's desire to not hold extreme risk-attitudes (i.e., risk-averse subjects do not want to become more risk-averse and risk-seeking subjects do not want to become even more risk-seeking).

We believe our work has important domain-specific implications for inventory management. Previous experimental work on algorithmic decision support in inventory management has primarily focused on the provision of risk-neutral order recommendations in simple single-item newsvendor tasks (Feng & Gao, 2020; Lee & Siemsen, 2017; Zhang & Siemsen, 2019). Such recommendations have been shown to minimize the pull-to-center bias, which is the tendency of DMs to place orders between the expected profit maximizing order quantity and the mean demand (Schweitzer & Cachon, 2000). In the present paper, we extend this line of work by showing that risk-averse algorithms could be used as anchors in complex tasks such as the multi-item newsvendor context to shift risk attitudes in a desirable direction. In fact, such algorithms could be used as a temporary behavioral intervention to inculcate lasting changes in ordering behavior.

Our study also provides practical insights for designing algorithmic decision support in inventory management when the firm's objective is to employ conservative recommender systems. Relying on risk-averse order strategies can be useful, since it leads to more stable profits, which is especially advantageous in situations involving chances of a loss. Our research highlights the critical need for firms to carefully select the level of risk aversion in the algorithm. Specifically, while we observed significant changes in ordering behavior with a highly risk-averse algorithm, we do not recommend consistently increasing the level of risk aversion since our results suggest that trust and advice utilization decrease as the level of risk aversion increases. In fact, it is important to understand the baseline risk attitude of managers to provide customized algorithmic recommendations to effectively induce changes in ordering behavior. Lastly, the fact that our findings remain consistent regardless of factors such as *autonomy* and *source of advice* is particularly encouraging for firms contemplating a swift adoption of inventory management software.

We now address some of the study's limitations. First, we focused on post-treatment effects, specifically the impact of algorithmic withdrawal, for a limited duration of twelve rounds. This constraint was necessary due to the nature of a lab study, where it is challenging to maintain participant engagement over extended periods. However, future research should examine the long-term effects of behavioral modifications induced by risk-averse algorithms. Second, we chose a simple algorithm design based on the mean-variance framework. It is possible that the impact of the algorithm might differ depending on its complexity. This variation in influence could be a subject for future research.

In sum, considering the growing relevance of algorithms in everyday life, this study offers insights into the important role algorithmic recommendations can play in influencing decision-making under risk. Our research demonstrates that in inventory management tasks that involve critical trade-offs between risk and return, algorithms can serve as a useful tool for purposefully inducing lasting changes in order decisions.

Funding

This work was funded by Fundação para a Ciência e a Tecnologia (UIDB/00124/2020, UIDP/00124/2020 and Social Sciences DataLab - PINFRA/22209/2016), POR Lisboa and POR Norte (Social Sciences DataLab, PINFRA/22209/2016).

CRediT authorship contribution statement

Pranadharthiharan Narayanan: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Jeeva Somasundaram:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Matthias Seifert:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Conceptualization.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.ejor.2024.11.013.

References

Ahmed, S., Çakmak, U., & Shapiro, A. (2007). Coherent risk measures in inventory problems. *European Journal of Operational Research, 182*(1), 226–238. https://doi.org/10.1016/J.EJOR.2006.07.016

Alsabah, H., Capponi, A., Ruiz Lacedelli, O., & Stern, M. (2021). Robo-advising: Learning investors' risk preferences via portfolio choices. *Journal of Financial Econometrics, 19*(2), 369–392. https://doi.org/10.1093/JJFINEC/NBZ040

Becker, G. S., & Murphy, K. M. (1988). A theory of rational addiction. *Journal of Political Economy, 96*(4), 675–700. https://doi.org/10.1086/261558

Becker-Peth, M., Thonemann, U. W., & Gully, T. (2018). A note on the risk aversion of informed newsvendors. *Journal of the Operational Research Society, 69*(7), 1135–1145. https://doi.org/10.1080/01605682.2017.1390525

Beggs, A., & Graddy, K. (2009). Anchoring effects: Evidence from art auctions. *American Economic Review, 99*(3), 1027–1039. https://doi.org/10.1257/aer.99.3.1027

Caro, F., & de Tejada Cuenca, A. S. (2023). Believing in analytics: Managers' adherence to price recommendations from a DSS. *Manufacturing & Service Operations Management, 25*(2), 524. https://doi.org/10.1287/MSOM.2022.1166

Chapman, G. B., & Johnson, E. J. (2012). Incorporating the irrelevant: Anchors in judgments of belief and value. *Heuristics and Biases*, 120–138. https://doi.org/10.1017/cbo9780511808098.008

Chen, X., Sim, M., Simchi-Levi, D., & Sun, P. (2007). Risk aversion in inventory management. *Operations Research, 55*(5), 828–842. https://doi.org/10.1287/opre.1070.0429

Choi, S., & Ruszczyński, A. (2011). A multi-product risk-averse newsvendor with exponential utility function. *European Journal of Operational Research, 214*(1), 78–84. https://doi.org/10.1016/J.EJOR.2011.04.005

D'Acunto, F., Prabhala, N., & Rossi, A. G. (2019). The promises and pitfalls of robo-advising. *Review of Financial Studies, 32*(5), 1983–2020. https://doi.org/10.1093/rfs/hhz014

De Véricourt, F., Jain, K., Bearden, J. N., & Filipowicz, A. (2013). Sex, risk and the newsvendor. *Journal of Operations Management, 31*(1–2), 86–92. https://doi.org/10.1016/J.JOM.2012.11.001

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General, 144*(1), 114–126. https://doi.org/10.1037/xge0000033

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science, 64*(3), 1155–1170. https://doi.org/10.1287/mnsc.2016.2643

Englich, B., & Mussweiler, T. (2001). Sentencing under uncertainty: Anchoring effects in the courtroom. *Journal of Applied Social Psychology, 31*(7), 1535–1551. https://doi.org/10.1111/j.1559-1816.2001.tb02687.x

Englich, B., Mussweiler, T., & Strack, F. (2005). The last word in court-a hidden disadvantage for the defense. *Law and Human Behavior, 29*(6). https://doi.org/10.1007/s10979-005-8380-7

Erlebacher, S. J. (2000). Optimal and heuristic solutions for the multi-item newsvendor problem with a single capacity constraint. *Production and Operations Management, 9*(3), 303–318. https://doi.org/10.1111/j.1937-5956.2000.tb00139.x

Feng, X., & Gao, J. (2020). Is optimal recommendation the best? A laboratory investigation under the newsvendor problem. *Decision Support Systems, 131*, Article 113251. https://doi.org/10.1016/j.dss.2020.113251

Fildes, R., Goodwin, P., Lawrence, M., & Nikolopoulos, K. (2009). Effective forecasting and judgmental adjustments: An empirical evaluation and strategies for improvement in supply-chain planning. *International Journal of Forecasting, 25*(1), 3–23. https://doi.org/10.1016/j.ijforecast.2008.11.010

Föbus, N., Lange, T., Leopoldseder, M., & Magnus, K.-H. (2019). *Food retail: The path to autonomous planning.* McKinsey Accessed Nov 11, 2024. Retrieved from https://www.mckinsey.com/industries/retail/our-insights/the-invisible-hand-on-the-path-to-autonomous-planning-in-food-retail.

Fujiwara, I., Ichiue, H., Nakazono, Y., & Shigemi, Y. (2013). Financial markets forecasts revisited: Are they rational, stubborn or jumpy? *Economics Letters, 118*(3), 526–530. https://doi.org/10.1016/J.ECONLET.2012.12.037

Glöckner, A., & Englich, B. (2015). When relevance matters: Anchoring effects can be larger for relevant than for irrelevant anchors. *Social Psychology*. https://doi.org/10.1027/1864-9335/a000214

Himmelstein, M., & Budescu, D. V. (2023). Preference for human or algorithmic forecasting advice does not predict if and how it is used. *Journal of Behavioral Decision Making, 36*(1), e2285. https://doi.org/10.1002/BDM.2285

Kesavan, S., & Kushwaha, T. (2020). Field experiment on the profit implications of merchants' discretionary power to override data-driven decision-making tools. *Management Science, 66*(11), 5182–5190. https://doi.org/10.1287/mnsc.2020.3743

Khosrowabadi, N., Hoberg, K., & Imdahl, C. (2022). Evaluating human behaviour in response to AI recommendations for judgemental forecasting. *European Journal of Operational Research, 303*(3), 1151–1167. https://doi.org/10.1016/J.EJOR.2022.03.017

Koehler, D. J., & Beauregard, T. A. (2006). Illusion of confirmation from exposure to another's hypothesis. *Journal of Behavioral Decision Making, 19*(1), 61–78. https://doi.org/10.1002/bdm.513

Lakens, D. (2024). When and How to Deviate From a Preregistration. *Collabra: Psychology, 10*(1), 2024. https://doi.org/10.1525/COLLABRA.117094/200749

Lee, Y. S., & Siemsen, E. (2017). Task decomposition and newsvendor decision making. *Management Science, 63*(10), 3226–3245. https://doi.org/10.1287/mnsc.2016.2521

Lieder, F., Griffiths, T. L., & Goodman, N. D. (2012). Burn-in, bias, and the rationality of anchoring. *Advances in Neural Information Processing Systems, 4*, 2690–2698.

Lin, W., Kim, S.-H., & Tong, J. (2023). What drives algorithm use? An empirical analysis of algorithm use in type 1 diabetes self-management. *SSRN Electronic Journal*. https://doi.org/10.2139/SSRN.3891832

Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes, 151*, 90–103. https://doi.org/10.1016/j.obhdp.2018.12.005

Loschelder, D. D., Trötschel, R., Swaab, R. I., Friese, M., & Galinsky, A. D. (2016). The information-anchoring model of first offers: When moving first helps versus hurts negotiators. *Journal of Applied Psychology, 101*(7), 995–1012. https://doi.org/10.1037/apl0000096

McAlvanah, P., & Moul, C. C. (2013). The house doesn't always win: Evidence of anchoring among Australian bookies. *Journal of Economic Behavior and Organization, 90*, 87–99. https://doi.org/10.1016/j.jebo.2013.03.009

Niszczota, P., & Kaszás, D. (2020). Robo-investment aversion. *PloS one, 15*(9), Article e0239277. https://doi.org/10.1371/JOURNAL.PONE.0239277

Prakash Katariya, A., Cetinkaya, S., Tekin, E., & C¸etinkaya, S. (2014). On the comparison of risk-neutral and risk-averse newsvendor problems. *Journal of the Operational Research Society, 65*(7), 1090–1107. https://doi.org/10.1057/jors.2013.48

Rader, C. A., Larrick, R. P., & Soll, J. B. (2017). Advice as a form of social influence: Informational motives and the consequences for accuracy. *Social and Personality Psychology Compass, 11*(8), e12329. https://doi.org/10.1111/SPC3.12329

Rader, C. A., Soll, J. B., & Larrick, R. P. (2015). Pushing away from representative advice: Advice taking, anchoring, and adjustment. *Organizational Behavior and Human Decision Processes, 130*, 26–43. https://doi.org/10.1016/j.obhdp.2015.05.004

Rebholz, T. R., Biella, M., & Hütter, M. (2024). Mixed-effects regression weights for advice taking and related phenomena of information sampling and utilization. *Journal of Behavioral Decision Making, 37*(2), e2369.

Sargent, T. J. (2009). *Dynamic macroeconomic theory*. Harvard University Press.

Schweitzer, M. E., & Cachon, G. P. (2000). Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science, 46*(3), 404–420. https://doi.org/10.1287/mnsc.46.3.404.12070

Sniezek, J. A., & Buckley, T. (1995). Cueing and cognitive conflict in judge-advisor decision making. *Organizational Behavior and Human Decision Processes, 62*(2), 159–174. https://doi.org/10.1006/obhd.1995.1040

Soll, J. B., & Larrick, R. P. (2009). Strategies for Revising Judgment: How (and How Well) People Use Others' Opinions. *Journal of Experimental Psychology: Learning Memory and Cognition, 35*(3), 780–805. https://doi.org/10.1037/A0015145

Soll, J. B., Palley, A. B., & Rader, C. A. (2022). The bad thing about good advice: Understanding when and how advice exacerbates overconfidence. *Management Science, 68*(4), 2949–2969. https://doi.org/10.1287/mnsc.2021.3987

Turner, B. M., & Schley, D. R. (2016). The anchor integration model: A descriptive model of anchoring effects. *Cognitive Psychology, 90*, 1–47. https://doi.org/10.1016/j.cogpsych.2016.07.003

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science (New York, N.Y.), 185*(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

Uppari, B. S., & Hasija, S. (2019). Modeling newsvendor behavior: A prospect theory approach. *Manufacturing and Service Operations Management, 21*(3), 481–500. https://doi.org/10.1287/msom.2017.0701

van de Kuilen, G., & Wakker, P. P. (2011). The midweight method to measure attitudes toward risk and ambiguity. *Management Science, 57*(3). https://doi.org/10.1287/mnsc.1100.1282

Wang, C. X., Webster, S., & Suresh, N. C. (2009). Would a risk-averse newsvendor order less at a higher selling price? *European Journal of Operational Research, 196*(2), 544–553. https://doi.org/10.1016/J.EJOR.2008.04.002

Wegener, D. T., Petty, R. E., Blankenship, K. L., & Detweiler-Bedell, B. (2010). Elaboration and numerical anchoring: Implications of attitude theories for consumer judgment and decision making. *Journal of Consumer Psychology, 20*(1), 5–16. https://doi.org/10.1016/j.jcps.2009.12.003

Wu, J., Li, J., Wang, S., & Cheng, T. C. E. (2009). Mean-variance analysis of the newsvendor model with stockout cost. *Omega, 37*, 724–730. https://doi.org/10.1016/j.omega.2008.02.005

Yaniv, I., & Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes, 83*(2), 260–281. https://doi.org/10.1006/obhd.2000.2909

Yoon, S., & Fong, N. (2019). Uninformative anchors have persistent effects on valuation judgments. *Journal of Consumer Psychology, 29*(3), 391–410. https://doi.org/10.1002/jcpy.1091

Zhang, Y., & Siemsen, E. (2019). A meta-analysis of newsvendor experiments: Revisiting the pull-to-center asymmetry. *Production and Operations Management, 28*(1), 140–156. https://doi.org/10.1111/poms.12899