



IE UNIVERSIDAD

TESIS DOCTORAL/  
DOCTORAL DISSERTATION

Soporte algorítmico con aversión al riesgo en problemas  
de asignación de recursos

Risk-averse algorithmic support in resource allocation  
problems

Pranadharthiharan Narayanan

SEGOVIA, 2024



IE UNIVERSIDAD

TESIS DOCTORAL/  
DOCTORAL DISSERTATION

Soporte algorítmico con aversión al riesgo en problemas  
de asignación de recursos

Risk-averse algorithmic support in resource allocation  
problems

Pranadharthiharan Narayanan

Doctoral Thesis Co-Advisors:  
Professor Jeeva Somasundaram  
Professor Matthias Seifert

## Abstract

We live in the age of algorithms. In this dissertation, I study how temporary exposure to risk-averse algorithms can influence decision-making in classic resource allocation problems such as portfolio management in finance and inventory management in operations. Using the anchoring and adjustment heuristic, I derive my predictions regarding algorithmic influence and test them using laboratory experiments. In chapters 1 and 2, I focus on project portfolio management and the multi-item newsvendor problem, respectively. In both these domains, I find that highly risk-averse algorithmic recommendations have the strongest influence on resource allocation decisions, despite individuals hedging away from the advice. Importantly, the changes in resource allocation decisions tend to persist even after the algorithm is no longer available. Chapter 3 reveals that these effects are similar regardless of factors such as decision autonomy (i.e., whether the algorithm is externally assigned or chosen by the subjects themselves) and source of advice (i.e., human or algorithm). Additionally, I find risk-averse algorithms can be used to counteract the “pull-to-center” bias in the low-profit newsvendor regime. Overall, I demonstrate the mutability of human behavior when temporarily exposed to risk-averse algorithmic aids. The findings are of notable consequence to firms strategically looking to utilize risk-averse algorithmic tools to improve resource allocation decisions without curtailing managerial autonomy.

## Resumen

Vivimos en la era de los algoritmos. En esta disertación, estudio cómo la exposición temporal a algoritmos adversos al riesgo puede influir en la toma de decisiones en problemas clásicos de asignación de recursos, como la gestión de carteras en finanzas y la gestión de inventarios en operaciones. Utilizando la heurística de anclaje y ajuste, obtengo mis predicciones sobre la influencia algorítmica y las pruebo mediante experimentos de laboratorio. En los capítulos 1 y 2, me centro en la gestión de la cartera de proyectos y el problema del vendedor de periódicos con artículos múltiples, respectivamente. En estos dominios, encuentro que las recomendaciones algorítmicas altamente reacias al riesgo tienen la mayor influencia en las decisiones de asignación de recursos, a pesar de que las personas evitan el consejo. Es importante destacar que los cambios en las decisiones de asignación de recursos tienden a persistir incluso después de que el algoritmo ya no esté disponible. El Capítulo 3 revela que estos efectos son similares independientemente de factores como la autonomía de decisión (es decir, si el algoritmo es asignado externamente o elegido por los propios sujetos) y la fuente de asesoramiento (es decir, humano o algoritmo). Además, encuentro que se pueden utilizar algoritmos con aversión al riesgo para mitigar significativamente el sesgo de “atracción hacia el centro” en el régimen de vendedores de periódicos de bajos beneficios. En general, demuestro la mutabilidad del comportamiento humano cuando se exponen temporalmente a ayudas algorítmicas con aversión al riesgo. Los hallazgos tienen consecuencias notables para las empresas que buscan estratégicamente utilizar herramientas algorítmicas adversas al riesgo para mejorar las decisiones de asignación de recursos sin restringir la autonomía gerencial.

## Acknowledgment

“What am I then...? Everything that I have seen, heard, and observed I have collected and exploited. My works have been nourished by countless different individuals, by innocent and wise ones, people of intelligence and dunces. Childhood, maturity, and old age all have brought me their thoughts..., their perspectives on life. I have often reaped what others have sowed. My work is the work of a collective being...” - Goethe

This is the part of the dissertation I *most* wanted to write. First, I want to thank my advisers Jeeva Somasundaram and Matthias Seifert for their guidance and friendship. Anything good that you read in these pages is thanks to their training and encouragement. I want to thank Kriti Jain for the gift of second chances. My heartfelt gratitude to my other committee members Emmanouil Avgerinos, Shameek Sinha, Konstantinos Stouras and Hamid Bouchikhi for their time and patience in providing valuable feedback. My sincere thanks to Laura Maguire and Maria Muriel for their unwavering support and reassuring presence since the beginning of the doctoral program.

I’m enormously grateful and ridiculously lucky for the unconditional love and support of my family (Amma, Appa, Anu and Maha). Their sacrifices have paved the way for my freedom. Finally, a big thank you to my wonderful friends Sreyaa, Sumit, Douglas, Shijith, Raj and Harsh who made this experience wholesome and memorable.

*Sarvam Krishnarpanam astu (Everything I offer to Krishna)*

# List of Contents

<b>Introduction .....</b>	<b>1</b>
<b>Introducción .....</b>	<b>7</b>
<b>Chapter 1: Risk-averse algorithmic support in project portfolio management.....</b>	<b>14</b>
1. Introduction.....	15
2. Theoretical framework .....	18
2.1 Anchoring and adjustment.....	18
2.2 Resource allocation problem .....	19
3. Experimental Study .....	22
3.1 Methods.....	22
3.2 Results .....	25
4. Discussion .....	29
<b>Chapter 2: Risk-averse algorithmic support in inventory management.....</b>	<b>32</b>
1. Introduction.....	33
2. Multi-item newsvendor model.....	36
3. Experimental Study .....	38
3.1 Methods.....	38
3.2 Results .....	41
3.3 Persistence of treatment effects. ....	45
4. Discussion .....	47
<b>Chapter 3: Risk-averse algorithmic support in inventory management: Further Extensions .....</b>	<b>49</b>
1. Introduction.....	50
2. Single-item newsvendor model.....	52
3. Study 1 .....	55
3.1 Methods.....	55
3.2 Results .....	58
4. Study 2 .....	62
4.1 Study 2a .....	63
4.2 Methods.....	64
4.3 Results .....	64
4.4 Study 2b .....	66

4.5 Methods.....	66
4.6 Results. ....	67
5. Discussion .....	68
<b>Conclusions.....</b>	<b>70</b>
<b>Conclusiones.....</b>	<b>74</b>
Appendix A.....	79
Appendix B.....	95
Appendix C.....	111
<b>References.....</b>	<b>136</b>

## List of Figures

Figure 1. Thesis overview .....	6
Figura 1. Visión general de la tesis .....	13
Figure 2. Study design .....	24
Figure 3. Endowment Allocated to MR Project.....	27
Figure 4. Study Design.....	40
Figure 5. Allocation to Product B.....	42
Figure 6. Average Standard Deviation in Profits .....	44
Figure 7. Study Design.....	57
Figure 8. (Mean) Units of Fish Ordered.....	59
Figure 9. Average Standard Deviation in Profits .....	61
Figure 10. (Mean) Pull to Center Bias .....	61
Figure 11. Allocation to Product B (Endogenous v. Exogenous Conditions) .....	65
Figure 12. Allocation to Product B.....	67
Figure 13. Past projects´ returns shown to participants in the study .....	89
Figure 14. Feedback screen displayed to a hypothetical participant in the study .....	89
Figure 15. Allocation to Product B v. Risk Aversion ( $\lambda$ ) .....	95
Figure 16. Feedback screen displayed to a hypothetical participant .....	106
Figure 17. Risk-averse algorithmic suggestion shown to participants .....	118
Figure 18. Feedback screen displayed to a hypothetical participant .....	118
Figure 19. Mean deviation from algorithmic recommendation.....	119
Figure 20. Balance of covariates before and after matching in study 1 .....	120
Figure 21. PtC bias - Within-subject analysis .....	122
Figure 22. PtC Score Distribution .....	122
Figure 23. Stimuli shown to participants prior to start of treatment period.....	126
Figure 24. Balance of covariates before and after matching .....	129
Figure 25. WOA Distribution.....	129
Figure 26. Study 2b design .....	131
Figure 27. Sample stimuli used in study 2b.....	132
Figure 28. Recommendation preference .....	134
Figure 29. Balance of covariates before and after matching in study 2b.....	134
Figure 30. WOA Distribution.....	135

## List of Tables

Table 1. OLS regression on resource allocation decisions with round-fixed effects .	28
Table 2. Demand distribution of products A and B .....	39
Table 3. OLS regression on ordering behavior with round-fixed effects .....	43
Table 4. IV regression on mean change in allocation to product B .....	46
Table 5. OLS regression on ordering behavior with round fixed-effects .....	59
Table 6. OLS regression on risk taking behavior with round-fixed effects.....	91
Table 7. OLS regression on deviation from algorithm with round-fixed effects.....	92
Table 8. OLS regression on deviation from algorithm with round-fixed effects.....	93
Table 9. Returns earned in the study .....	94
Table 10. OLS regression on (absolute) deviation from algorithmic recommendation with round-fixed effects. ....	107
Table 11. OLS regression on (normalized) deviation from algorithmic recommendation with round-fixed effects.....	108
Table 12. Profits earned .....	109
Table 13. Alternate explanations for study results.....	109
Table 14. OLS regression on profits earned with round-fixed effects.....	121
Table 15. Demand chasing analysis in study 1 .....	124
Table 16. OLS regression on ordering behaviour with round-fixed effects .....	127
Table 17. OLS regression on weight of advice with round-fixed effects .....	130
Table 18. OLS regression ordering behavior with round-fixed effects .....	133
Table 19. OLS regression on weight of advice with round-fixed effects.....	135

## Introduction

We live in the age of algorithms. They feel omnipresent in remarkable ways, ranging from the mundane (e.g., finding a cab) to the extraordinary (e.g., DNA sequencing). The rise of small devices and big data in conjunction with the development of machine learning tools and artificial intelligence have squarely placed algorithms at the focus of data-driven decision-making. It is in this backdrop that I write this dissertation, wherein I focus on the *human-algorithm* connection, a topic of great interest to a diverse community of scholars in decision sciences, behavioral operations and various others sub-fields of management. The broad agenda for this dissertation, as also stated in the latest call for papers in the special issue of *Management Science*, is to help humans co-exist better with the algorithms that surround us, by understanding the (subtle) ways in which algorithms affect human behavior (Caro et al., 2021).

My research primarily investigates how algorithms influence decision-making under risk, particularly in the context of allocating resources. I primarily focus on a class of resource allocation problems where a decision-maker (DM) must distribute a fixed or constrained resource, such as money, between two options where one alternative is riskier than the other (e.g., stocks v. bonds). By 'risk', I mean the variability in returns, quantified as the standard deviation-to-mean ratio (i.e., coefficient of variation). This definition aligns with the economic perspective of risk, differing from its common interpretation as the likelihood of adverse outcomes (e.g., the risk of contracting a disease). Furthermore, when I mention 'algorithm', I'm referring to a recommendation of optimal allocation of resources, like investing 60% in stocks and 40% in bonds, which is derived from solving an underlying optimization problem. In

my dissertation, I interchangeably use the terms 'algorithm' and 'algorithmic recommendation'.

Algorithms are increasingly used in diverse resource allocation tasks such as inventory management in operations (Bertsimas & Kallus, 2016) and portfolio management in finance (D'Acunto et al., 2019). For instance, in the domain of online financial services, “robo-advisors” are becoming increasingly popular worldwide as investors are seeking automated and cost-effective investment solutions (Source: [Statista](#)). In the retail sector, algorithms are now routinely used for tasks such as demand forecasting, pricing, and inventory replenishment (Kesavan & Kushwaha, 2020a). The use of algorithmic recommendations helps firms improve operational efficiency, reduce costs and potentially counteract human biases in decision-making (Feng & Gao, 2020; Lee & Siemsen, 2017; Schweitzer & Cachon, 2000). Despite the proliferation of algorithms, 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 (see, for instance, Lin et al., 2023). This dissertation extends this line of research by studying behavioral responses to algorithms that are calibrated with specific, pre-determined levels of risk aversion in the context of resource allocation problems. In other words, *how do people with different risk attitudes respond to algorithms with different risk aversion levels in resource allocation problems?* Addressing this question is at the heart of this dissertation.

Risk preferences play a fundamental role in many resource allocation problems such as portfolio management in finance (Markowitz, 1952) and inventory

management in operations (Chen et al., 2007). For instance, in portfolio optimization, risk-averse algorithms which are customized according to a decision-makers' risk attitude are commonly used to aid investment decisions (Alsabah et al., 2021; Bhattacharya et al., 2012; D'Acunto et al., 2019). Even in inventory management, 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 DMs in a risk-averse direction. Importantly, risk-averse decisions help reduce the volatility in profits or returns. However, I am not aware of any behavioral work that explicitly focuses on responses to well-calibrated algorithmic advice of consistent level(s) of risk-aversion in resource allocation problems.

Additionally, extant studies on algorithmic decision support in these contexts do not examine if algorithmic recommendations can influence behavior even *after the algorithm is no longer available*. In a world where practitioners increasingly resort to algorithmic recommendations (i.e., software solutions) that are deployed via subscription-based licensing agreements, it is crucial for companies to assess whether any behavioral modifications induced by algorithms are likely to persist even after the expiry of the subscription period (i.e., once the algorithmic support is withdrawn). Therefore, understanding the effects of *temporary* risk-averse algorithmic support is another important agenda of this dissertation.

In Chapters 1 and 2, I explore how algorithms programmed with specific, predetermined levels of risk aversion can influence DMs in resource allocation tasks such as project portfolio management (chapter 1) and inventory management (chapter

2). I derive my predictions using the anchoring and adjustment heuristic (Tversky & Kahneman, 1974) and test the predictions in the lab across two diverse resource problems: project portfolio management (chapter 1) and multi-item newsvendor task (chapter 2). Interestingly – for this first time – I show that highly risk-averse (HRA) algorithms have the strongest effect on resource allocation decisions, despite individuals relying on HRA algorithms the least and hedging away from the advice. Importantly, the change in resource allocation decisions (in a *more* risk-averse direction) persists even after the (HRA) algorithm is removed. That is, people act risk-averse after being temporarily exposed to HRA algorithms across the two diverse resource allocation contexts. The fact that risk-adverse algorithmic aids can have an enduring effect on human behavior has important policy and managerial implications. For example, companies might want to use algorithmic support to nudge managers to maintain a certain level of risk that is aligned with the firm’s strategic goals.

In chapter 3, I explore the boundary conditions of the observed algorithmic influence in inventory management tasks. Specifically, I examine whether critical factors, such as decision autonomy (i.e., choosing an algorithm of a particular risk-level versus random assignment) and source of advice (human versus algorithm) have an impact on the extent of algorithmic utilization. This investigation is motivated from prior research on algorithm aversion and advice-taking, which suggest that there could be differences in algorithmic use depending on whether the algorithmic assignment is endogenous or exogenous (Dietvorst et al., 2018; Rader et al., 2017) and whether the recommendations are from an algorithm or a human (Dietvorst et al., 2015). My main finding is that anchoring effects (i.e., algorithmic influence) are similar regardless of factors such as decision autonomy and source of advice. Finally, I also demonstrate

that risk-averse algorithms can be a potential debiasing tool in the classic, single-item newsvendor task.

Throughout this dissertation, laboratory studies serve as the primary method for my research. The experimental design draws inspiration from the field of habit formation (Becker & Murphy, 1988). This approach involves not only observing how individuals respond to algorithms (during a treatment period) but also assessing their reactions after the algorithmic aids are removed (in a post-treatment period). This distinctive aspect of the research design enables both between-subjects and within-subject comparisons, thereby enhancing the reliability of the findings.

The results from this dissertation have important implications for the academic literature on decision-making and operations management. First, I contribute to the extant work on anchoring by showing that risk-averse algorithms, which contain informational value, can leave an imprint on the DM following temporary exposure. This is unlike recent studies that have demonstrated the persistence of anchoring only using “uninformative” anchors (Yoon & Fong, 2019). Second, prior research on decision support in inventory management has predominantly focused on risk-neutral recommendations in simple, single-item newsvendor tasks (Feng & Gao, 2020b; Lee & Siemsen, 2017; Zhang & Siemsen, 2019). In this dissertation, I extend this line of work by showing that risk-averse algorithms could be used as anchors in more complex inventory tasks such as the multi-item newsvendor problem to inculcate lasting changes in ordering behavior. Beyond contributions to the academic literature, the findings have important, practical domain-specific implications in resource allocation problems. For instance, in the domain of finance, subscription-based decision support is becoming increasingly common not only for professional traders, but also for lay people investing their own private wealth. My results show that in these

contexts, companies could use algorithmic recommendations to nudge employees or consumers towards a specific risk-taking behavior. A brief overview of the thesis is shown in Figure 1 below.

Chapter 1	Chapter 2	Chapter 3
<p><b>Key research question(s):</b></p>		
<p>What is the effect of algorithm(s) of different risk aversion levels on investment decisions in project portfolio management?</p>	<p>What is the effect of algorithm(s) of different risk aversion levels on ordering decisions in inventory management?</p>	<p>How does algorithmic influence vary in inventory management depending on factors such as decision autonomy (endogenous v. exogenous choice) and source of advice (human v. algorithm)?</p>
<p><b>Phenomena of Interest:</b></p>		
<p>Persistence of anchoring (i.e., stickiness of algorithmic recommendations)</p>	<p>Persistence of anchoring (i.e., stickiness of algorithmic recommendations)</p>	<p>Algorithm aversion Pull-to-Center bias</p>
<p><b>Main Finding(s):</b></p>		
<p>Subjects become more risk averse when temporarily exposed to highly risk averse algorithms, despite hedging away from it.</p>	<p>Subjects become more risk averse when temporarily exposed to highly risk averse algorithms, despite hedging away from it.</p>	<p>No difference in advice utilization based on either decision autonomy or source of advice.</p> <p>An important implication in the single-item newsvendor task is that risk-averse algorithms can be used to minimize the pull-to-center bias.</p>

Figure 1. Thesis overview

## Introducción

Vivimos en la era de los algoritmos. Se sienten omnipresentes de maneras notables, que van desde lo mundano (por ejemplo, encontrar un taxi) hasta lo extraordinario (por ejemplo, la secuenciación del ADN). El auge de los dispositivos pequeños y los grandes datos, en conjunto con el desarrollo de herramientas de aprendizaje automático e inteligencia artificial, han colocado a los algoritmos en el centro de la toma de decisiones basada en datos. Es en este contexto que escribo esta disertación, en la que me enfoco en la conexión humano-algoritmo, un tema de gran interés para una comunidad diversa de académicos en ciencias de la decisión, operaciones conductuales y varios otros subcampos de la gestión. La agenda amplia de esta disertación, como también se menciona en la última convocatoria de trabajos en el número especial de *Management Science*, es ayudar a los humanos a coexistir mejor con los algoritmos que nos rodean, comprendiendo las formas (sutiles) en que los algoritmos afectan el comportamiento humano (Caro et al., 2021).

Mi investigación investiga principalmente cómo los algoritmos influyen en la toma de decisiones bajo riesgo, particularmente en el contexto de la asignación de recursos. Me enfoco principalmente en una clase de problemas de asignación de recursos donde un tomador de decisiones (DM) debe distribuir un recurso fijo o limitado, como dinero, entre dos opciones donde una alternativa es más riesgosa que la otra (por ejemplo, acciones versus bonos). Por 'riesgo', me refiero a la variabilidad en los retornos, cuantificada como la relación entre la desviación estándar y la media (es decir, coeficiente de variación). Esta definición se alinea con la perspectiva económica del riesgo, diferenciándose de su interpretación común como la probabilidad de resultados adversos (por ejemplo, el riesgo de contraer una enfermedad). Además, cuando menciono 'algoritmo', me refiero a una recomendación

de asignación óptima de recursos, como invertir el 60% en acciones y el 40% en bonos, que se deriva de resolver un problema de optimización subyacente. En mi disertación, uso indistintamente los términos 'algoritmo' y 'recomendación algorítmica'. De manera similar, los pronombres "Yo" y "Nosotros" se usan de manera sinónima a lo largo del documento.

Los algoritmos se utilizan cada vez más en diversas tareas de asignación de recursos, como la gestión de inventarios en operaciones (Bertsimas & Kallus, 2016) y la gestión de carteras en finanzas (D'Acunto et al., 2019). Por ejemplo, en el dominio de los servicios financieros en línea, los “robo-asesores” se están volviendo cada vez más populares en todo el mundo, ya que los inversores buscan soluciones de inversión automatizadas y rentables (Fuente: [Statista](#)). En el sector minorista, los algoritmos se utilizan ahora rutinariamente para tareas como pronóstico de demanda, precios y reposición de inventarios (Kesavan & Kushwaha, 2020a). El uso de recomendaciones algorítmicas ayuda a las empresas a mejorar la eficiencia operativa, reducir costos y potencialmente contrarrestar los sesgos humanos en la toma de decisiones (Feng & Gao, 2020; Lee & Siemsen, 2017; Schweitzer & Cachon, 2000). A pesar de la proliferación de algoritmos, investigaciones recientes muestran que las personas tienden a hacer ajustes de juicio a las recomendaciones algorítmicas o anularlas por completo, incluso si esto conlleva un rendimiento más bajo (Dietvorst et al., 2015; Fildes et al., 2009; Khosrowabadi et al., 2022). Estudiar esta tendencia conductual y encontrar formas de superar lo que ahora se conoce como “aversión al algoritmo” se ha convertido en una importante agenda de investigación (ver, por ejemplo, Lin et al., 2023). Esta disertación extiende esta línea de investigación al estudiar las respuestas conductuales a algoritmos que están calibrados con niveles específicos y predefinidos de aversión al riesgo en el contexto de problemas de

asignación de recursos. En otras palabras, ¿cómo responden las personas con diferentes actitudes hacia el riesgo a algoritmos con distintos niveles de aversión al riesgo en problemas de asignación de recursos? Abordar esta pregunta es el núcleo de esta disertación.

Las preferencias de riesgo juegan un papel fundamental en muchos problemas de asignación de recursos, como la gestión de carteras en finanzas (Markowitz, 1952) y la gestión de inventarios en operaciones (Chen et al., 2007). Por ejemplo, en la optimización de carteras, se utilizan comúnmente algoritmos reacios al riesgo que están personalizados según la actitud de riesgo de los tomadores de decisiones para ayudar en las decisiones de inversión (Alsabah et al., 2021; Bhattacharya et al., 2012; D'Acunto et al., 2019). Incluso en la gestión de inventarios, en muchos escenarios que involucran costos de almacenamiento significativos (por ejemplo, suministros médicos sensibles al tiempo), ciclos de vida cortos del producto (por ejemplo, electrónica de consumo), alta incertidumbre de la demanda (por ejemplo, venta al por menor de moda), restricciones de capital (por ejemplo, pequeñas empresas) o bienes perecederos sin valor de rescate (por ejemplo, productos frescos), puede ser deseable utilizar consejos algorítmicos para influir en las preferencias de riesgo (desconocidas) de los DM en una dirección reacia al riesgo. Importante, las decisiones reacias al riesgo ayudan a reducir la volatilidad en las ganancias o retornos. Sin embargo, no conozco ningún trabajo conductual que se enfoque explícitamente en las respuestas a consejos algorítmicos bien calibrados de nivel(es) consistente(s) de aversión al riesgo en problemas de asignación de recursos.

Además, los estudios existentes sobre apoyo a la decisión algorítmica en estos contextos no examinan si las recomendaciones algorítmicas pueden influir en el comportamiento incluso después de que el algoritmo ya no esté disponible. En un

mundo donde los profesionales recurren cada vez más a recomendaciones algorítmicas (es decir, soluciones de software) que se implementan a través de acuerdos de licencias basados en suscripciones, es crucial para las empresas evaluar si las modificaciones conductuales inducidas por los algoritmos probablemente persistirán incluso después del vencimiento del período de suscripción (es decir, una vez que se retire el apoyo algorítmico). Por lo tanto, comprender los efectos del apoyo algorítmico temporal reacio al riesgo es otra agenda importante de esta disertación.

En los capítulos 1 y 2, exploro cómo los algoritmos programados con niveles específicos y predefinidos de aversión al riesgo pueden influir en los tomadores de decisiones (DMs) en tareas de asignación de recursos, como la gestión de carteras de proyectos (capítulo 1) y la gestión de inventarios (capítulo 2). Derivo mis predicciones utilizando la heurística de anclaje y ajuste (Tversky & Kahneman, 1974) y pruebo las predicciones en el laboratorio a través de dos problemas de recursos diversos: gestión de carteras de proyectos (capítulo 1) y tarea de nuevo proveedor multi-item (capítulo 2). Curiosamente, por primera vez, mostramos que los algoritmos altamente reacios al riesgo (HRA) tienen el efecto más fuerte en las decisiones de asignación de recursos, a pesar de que los individuos dependen menos de los algoritmos HRA y se alejan de sus consejos. Importante, el cambio en las decisiones de asignación de recursos (en una dirección más reacia al riesgo) persiste incluso después de que se retire el algoritmo (HRA). Es decir, las personas actúan de manera reacia al riesgo después de estar expuestas temporalmente a algoritmos HRA en los dos contextos diversos de asignación de recursos. El hecho de que las ayudas algorítmicas reacias al riesgo puedan tener un efecto duradero en el comportamiento humano tiene importantes implicaciones políticas y gerenciales. Por ejemplo, las empresas podrían querer utilizar el apoyo algorítmico para impulsar a los gerentes a

mantener un cierto nivel de riesgo alineado con los objetivos estratégicos de la empresa.

En el capítulo 3, exploro las condiciones límite de la influencia algorítmica observada en tareas de gestión de inventarios. Específicamente, examino si factores críticos, como la autonomía de decisión (es decir, elegir un algoritmo de un nivel de riesgo particular versus asignación aleatoria) y la fuente del consejo (humano versus algoritmo) tienen un impacto en el grado de utilización algorítmica. Esta investigación está motivada por investigaciones previas sobre aversión al algoritmo y la toma de consejos, que sugieren que podría haber diferencias en el uso algorítmico dependiendo de si la asignación algorítmica es endógena o exógena (Dietvorst et al., 2018; Rader et al., 2017) y si las recomendaciones provienen de un algoritmo o un humano (Dietvorst et al., 2015). Mi principal hallazgo es que los efectos de anclaje (es decir, la influencia algorítmica) son similares independientemente de factores como la autonomía de decisión y la fuente del consejo. Finalmente, también demuestro que los algoritmos reacios al riesgo pueden ser una herramienta potencial para eliminar sesgos en la clásica tarea del nuevo proveedor de un solo artículo.

A lo largo de esta disertación, los estudios de laboratorio sirven como el método principal para mi investigación. El diseño experimental se inspira en el campo de la formación de hábitos (Becker & Murphy, 1988). Este enfoque implica no solo observar cómo los individuos responden a los algoritmos (durante un período de tratamiento) sino también evaluar sus reacciones después de que se retiren las ayudas algorítmicas (en un período posterior al tratamiento). Este aspecto distintivo de nuestro diseño de investigación permite comparaciones tanto entre sujetos como dentro de un mismo sujeto, mejorando así la fiabilidad de los hallazgos.

Los resultados de esta disertación tienen importantes implicaciones para la literatura académica sobre toma de decisiones y gestión de operaciones. Primero, contribuimos al trabajo existente sobre anclaje al mostrar que los algoritmos reacios al riesgo, que contienen valor informativo, pueden dejar una huella en el DM después de una exposición temporal. Esto es diferente de estudios recientes que han demostrado la persistencia del anclaje utilizando solo anclas “no informativas” (Yoon & Fong, 2019). Segundo, la investigación previa sobre apoyo a la decisión en gestión de inventarios se ha centrado predominantemente en recomendaciones neutrales al riesgo en tareas simples de nuevo proveedor de un solo artículo (Feng & Gao, 2020b; Lee & Siemsen, 2017; Zhang & Siemsen, 2019). En esta disertación, extendiendo esta línea de trabajo al mostrar que los algoritmos reacios al riesgo podrían utilizarse como anclas en tareas de inventario más complejas, como el problema del nuevo proveedor multi-item, para inculcar cambios duraderos en el comportamiento de pedidos. Más allá de las contribuciones a la literatura académica, los hallazgos tienen importantes implicaciones prácticas específicas del dominio en problemas de asignación de recursos. Por ejemplo, en el dominio de las finanzas, el apoyo a la decisión basado en suscripciones se está volviendo cada vez más común no solo para operadores profesionales, sino también para personas comunes que invierten su propio patrimonio privado. Nuestros resultados muestran que en estos contextos, las empresas podrían utilizar recomendaciones algorítmicas para impulsar a empleados o consumidores hacia un comportamiento específico de toma de riesgos. A continuación se muestra un breve resumen de la tesis en la Figura 1.

Capítulo 1	Capítulo 2	Capítulo 3
<p data-bbox="268 277 572 344"><b>Pregunta(s) clave de investigación:</b></p> <p data-bbox="256 360 624 517">¿Cuál es el efecto de los algoritmo(s) de diferentes niveles de aversión al riesgo en las decisiones de inversión en la gestión de carteras de proyectos?</p> <p data-bbox="268 584 572 629"><b>Fenómeno de Interés:</b></p> <p data-bbox="256 667 619 757">Persistencia del anclaje (es decir, la adherencia de las recomendaciones algorítmicas)</p> <p data-bbox="268 824 572 869"><b>Hallazgo(s) Principal(es):</b></p> <p data-bbox="256 907 619 1064">Los sujetos se vuelven más aversos al riesgo cuando están expuestos temporalmente a algoritmos altamente aversos al riesgo, a pesar de alejarse de este.</p>	<p data-bbox="667 360 1021 517">¿Cuál es el efecto de los algoritmo(s) de diferentes niveles de aversión al riesgo en las decisiones de pedido en la gestión de inventarios?</p> <p data-bbox="687 667 1050 757">Persistencia del anclaje (es decir, la adherencia de las recomendaciones algorítmicas)</p> <p data-bbox="687 907 1050 1064">Los sujetos se vuelven más aversos al riesgo cuando están expuestos temporalmente a algoritmos altamente aversos al riesgo, a pesar de alejarse de este.</p>	<p data-bbox="1110 360 1493 551">¿Cómo varía la influencia algorítmica en la gestión de inventarios dependiendo de factores como la autonomía de decisión (endógena vs. exógena) y la fuente del consejo (humano vs. algoritmo)?</p> <p data-bbox="1134 667 1442 745">Aversión al algoritmo Sesgo de Tirar-hacia-el-Centro</p> <p data-bbox="1126 907 1493 996">No hay diferencia en la utilización de consejos basada en la autonomía de decisión o la fuente del consejo.</p> <p data-bbox="1126 1021 1469 1207">Una implicación importante en la tarea de nuevo proveedor de un solo artículo es que los algoritmos aversos al riesgo pueden usarse para minimizar el sesgo de Tirar-hacia-el-Centro.</p>

Figura 2. Visión general de la tesis

# Chapter 1: Risk-averse algorithmic support in project portfolio management

**Abstract:** I examine how resource allocation decisions change when decision-makers (DMs) are temporarily exposed to algorithmic aids that are programmed with specific, predetermined levels of risk aversion. Using the anchoring and adjustment heuristic, I derive my predictions and test them in a project portfolio management task. I find that highly risk-averse algorithmic recommendations have the strongest influence on resource allocation decisions, despite individuals relying on it the least and hedging away from the advice. Importantly, the changes in resource allocation decisions are sticky even after the algorithm is no longer available. The findings demonstrate the mutability of human decisions under temporary algorithmic influence.

**Keywords:** algorithm | risk aversion | anchoring | resource allocation

## 1. Introduction

People are often reluctant to trust in the output of imperfect algorithmic models, even when their predictive performance is known to be superior (Dietvorst et al., 2015). This finding has sparked follow up research in a wide range of disciplines including organizational behavior (Logg et al., 2019; Newman et al., 2020), finance (D'Acunto et al., 2019), marketing (Castelo et al., 2019; Luo et al., 2019; Puntoni et al., 2021) and operations (Bastani et al., 2021; Kesavan & Kushwaha, 2020). Chapter 1 extends this line of research by studying behavioral responses to algorithms that are calibrated with specific, pre-determined levels of risk aversion in the context of project portfolio management decisions.

More generally, risk preferences play a fundamental role in many resource allocation problems such as portfolio management in finance (Markowitz, 1952) and inventory management in operations (Chen et al., 2007). Importantly, risk-averse resource allocation decisions help reduce the volatility in profits or returns. For instance, in portfolio optimization, risk-averse algorithms which are customized according to a decision-makers' risk attitude are commonly used to aid investment decisions (Alsabah et al., 2021; Bhattacharya et al., 2012; D'Acunto et al., 2019). However, I am not aware of any empirical work that explicitly focus on behavioral responses to well-calibrated algorithmic advice of consistent level(s) of risk-aversion.

Additionally, extant studies on decision support in resource allocation tasks do not examine if algorithmic recommendations can influence behavior even *after the algorithm is no longer available*. I believe studying the issue of *algorithmic withdrawal* is important for several reasons. First, from a decision theoretic standpoint, understanding algorithmic withdrawal allows us to gain a more comprehensive understanding of algorithmic influence on resource allocation decisions by considering

both the presence and subsequent absence of the decision support system. Second, in a world where practitioners increasingly resort to licensing third party, off-the-shelf analytics 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. Finally, if the behavioral modification induced by algorithmic support does have a lasting impact, then firms can strategically deploy algorithmic support only for a limited duration of time. Therefore, it becomes vital for both practitioners and researchers to examine the lasting effects of temporary algorithmic support.

Using the anchoring and adjustment heuristic I predict how DMs' portfolio allocation decisions would change in response to algorithmic recommendations. I then test the prediction in a classic project portfolio management task, where subjects had an opportunity to allocate a fixed resource into two different prospects: (i) a more-risky (MR) alternative that offers a higher expected return; and (ii) a less-risky (LR) alternative offering a lower expected return. I specifically choose the project portfolio management context since it is cognitively challenging, and the necessity to allocate fixed resources (i.e., endowment) among multiple alternatives (i.e., projects with different volatilities in returns) makes the concept of risk highly salient from the perspective of the DM. The study was comprised of three experimental periods: a baseline period, a treatment period, and a post-treatment period. Each period consisted of several decision rounds.

During the baseline period, I studied subjects' allocation decisions without algorithmic recommendations. Subjects were then assigned either to a control condition, in which they did not encounter any algorithm, or to one of three treatment conditions in which subjects 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, the suggested allocation differed based on the algorithm's assumed level of risk. The RN algorithm allocated resources that maximized expected returns. In contrast, both HRA and SRA algorithms allocated less resources to the MR alternative than the RN algorithm. Subjects in all treatment conditions were free to choose whether or not to follow the algorithmic recommendations. Across all conditions, I examined how subjects allocated their resources between the two alternatives during the treatment and post-treatment periods (after the algorithm was removed). The experimental design allowed us to conduct both between-subjects and within-subject comparisons. For instance, I can analyze the allocation decisions of subjects in the treatment condition(s) in comparison to those in the control group at any given time (between-subjects). Furthermore, I can assess the allocation decisions of subjects in the treatment condition(s) both before and after they encounter the algorithm (within-subject). This distinctive aspect of the research design enhances the reliability of the findings.

In the study, I found an interesting tension between changes in allocation decisions and algorithmic utilization, which is a proxy for underlying strength of anchoring. Specifically, I 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 and hedging away from it.

## 2. Theoretical framework

### 2.1 Anchoring and adjustment

Anchoring in the decision-making literature 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).<sup>1</sup>

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 & Beaugregard, 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 (Bhatia, 2012.; Lieder et al., 2012; Turner & Schley, 2016).

The anchoring paradigm I propose in this paper has two key features: First, to predict treatment period effects, I assume that DMs anchor their decisions based on the algorithm's recommendation and insufficiently adjust towards an optimal level that is based on their own baseline decisions (i.e., anchoring-and-adjustment heuristic). Second, I allow for persistence of anchoring in the post-treatment period, which is conditional on treatment effects. I first describe anchoring below in the context of a general resource allocation setting and incorporate multiple algorithms of varying risk aversion levels that are assigned exogenously.

---

<sup>1</sup> In behavioral operations, Schweitzer and Cachon (2000) used anchoring on mean demand as a possible explanation for the PtC bias in the single item newsvendor problem. This account has recently been corroborated from the perspective of prospect theory (Uppari & Hasija, 2019)

## 2.2 Resource allocation problem

**Setup.** Consider a DM facing the task of investing her endowment (resource)  $W$  in two prospects: a more-risky (MR) and a less-risky (LR) alternative. I assume that the mean returns ( $\mu_1$ ) and coefficient of variation ( $\sigma_1/\mu_1$ ) from alternative MR are higher than the respective mean returns ( $\mu_2$ ) and coefficient of variation ( $\sigma_2/\mu_2$ ) from alternative LR. In other words, alternative MR is comparatively the *high-risk, high-return* option. Now, let  $\alpha$  and  $(1 - \alpha)$  be the fraction of the endowment allocated to alternatives MR and LR by the DM, respectively. The expected value of the returns is given by  $E = \alpha\mu_1 + (1 - \alpha)\mu_2$  and the variance of the returns<sup>2</sup> (i.e., risk) is given by  $V = \alpha^2\sigma_1^2 + (1 - \alpha)^2\sigma_2^2$ .

Consistent with modern portfolio theory, I assume that the DM tries to maximize the expected returns  $E$  while trying to minimize the variance  $V$  subject to her degree of risk aversion  $\lambda$  (Markowitz, 1952b). Therefore, I can represent the objective function  $F$  of the DM as follows:

$$F = E - \lambda V \quad (1)$$

The intuition behind Eq. (1) is that an increase in risk-aversion  $\lambda$  is modeled by an increasing distaste for payoff variance. Therefore, *higher values of risk aversion  $\lambda$  should lead to a lower allocation of resources  $\alpha$  toward the more-risky (MR) alternative* (Eq. (1) actually results from an expected utility maximization. See Appendix A1 for the derivation).

**Predictions.** The objective is to examine how DMs modify their resource allocation decisions in response to the algorithmic recommendations and derive

---

<sup>2</sup> In presenting the formula for the variance, we assumed zero correlation between the returns from the MR and LR alternatives. If the returns are correlated with a value  $\rho$ , we have  $V = \alpha^2\sigma_1^2 + (1 - \alpha)^2\sigma_2^2 + 2\alpha(1 - \alpha)\rho\sigma_1\sigma_2$

predictions for the experiment. Without loss of generality, I focus on the three main periods of the experiment: (i) Baseline period ( $t = 0$ ) where the DM makes their decisions without any algorithmic aid, (ii) Treatment period ( $t = 1$ ), in which DMs (randomly) assigned to treatment conditions make resource allocation 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 resource allocation 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. I allow  $\alpha$  in the 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 allocation decision in each experimental time period. I assume that  $\alpha_t$  evolves as follows:

**Baseline period:** As subjects are perfectly randomized, I do not expect to see any difference in (average)  $\alpha$  across conditions at the end of the baseline period i.e.,  $\alpha_0^{HRA} = \alpha_0^{SRA} = \alpha_0^{RN} = \alpha_0^{Control}$  at the end of baseline period.

**Treatment period:** The DM is assumed to compute her optimal allocation  $\alpha^*$  towards the MR alternative based on her degree of risk aversion  $\lambda$  using Eq. (1). However, during the treatment period, the DM's allocation decision could be influenced by a contextual anchor, such as an algorithmic recommendation that they encounter<sup>3</sup>.

---

<sup>3</sup> The risk aversion parameter  $\lambda$  and allocation  $\alpha$  can also be affected by fluctuations in wealth (Brunnermeier & Nagel, 2008; Steffensen, 2011). However, as subjects in the experiment do not accumulate wealth over rounds and are incentivized based on returns earned in a randomly selected round, we abstract away from factors such as temporary wealth held in a given round that might affect  $\alpha$ .

In the presence of an algorithm, I assume that the DM's final allocation  $A^*$  towards the MR alternative is a weighted average of optimal allocation  $\alpha^*$  and the algorithmic recommendation  $x^j$  towards the MR alternative, where  $j = HRA, SRA$  or  $RN$  represents the condition that DM is assigned to. The algorithmic recommendations are of varying levels of risk aversion ( $x^{HRA} < x^{SRA} < x^{RN}$ ). Therefore, using the anchoring and adjustment heuristic, the DM's final allocation towards the MR alternative during the treatment period is given by:

$$A^* = \beta\alpha^* + (1 - \beta)x^j \quad (2)$$

In Eq. (2),  $\beta > 0$  captures the weight DMs give to optimal allocation  $\alpha^*$  and the adjustment they make away from the anchor  $x^j$ . As  $x^{HRA} < x^{SRA} < x^{RN}$ , it follows from Eq. (2) that DMs who observe the HRA algorithm allocate the least to the MR alternative and are consequently the most risk-averse during the treatment period (*ceteris paribus*). This leads us to the first hypothesis.

***H1:** DMs who observe highly risk-averse algorithms are more risk-averse in their allocations compared to those who do not observe any algorithms or those who observe slightly risk-averse or risk neutral algorithms.*

**Post-treatment period:** Note that in the post-treatment period no anchor (algorithmic recommendation) is provided to DMs. However, I 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 & Fong, 2019) find that uninformative external anchors (i.e., random values uncorrelated with the true

estimate) still have a persistent effect on valuation judgments, such as DMs' willingness-to-pay (WTP). While recent research suggests that relevant anchors (such as the algorithmic recommendation I use) are expected to have larger effects compared to irrelevant anchors, the results are still mixed, especially regarding low anchors (Glöckner & Englich, 2015). Thus, based on findings from prior literature, I expect anchoring effects to persist and hypothesize the following and test the predictions in study 1.

***H2:** DMs who observe highly risk-averse algorithms continue to be risk-averse in their allocations even after the algorithm is removed.*

I test the hypotheses in two resource allocation tasks involving project portfolio management (chapter 1) and newsvendor environment (in chapter 2).

### **3. Experimental Study**

The experimental study draws on a project portfolio management task, in which subjects were required to allocate resources between two projects that differed in their level of risk.

#### **3.1 Methods**

**Subjects.** Subjects were Master's students at a large private university. The study was conducted online, and subjects earned class credits in return for their participation. In addition to earning extra course credits, I selected 10 subjects at random and awarded them Amazon vouchers matching the earnings of one randomly chosen round. The 10 winning subjects earned vouchers worth € 22 on average. In total, 197 subjects (Female = 42%) with an average age of 25 years participated in the study.

**Design and Procedure.** Once subjects consented to take part in the study and were briefed about the incentives, they were asked to assume the role of a manager and allocate their company's financial resources to two projects: a more-risky project (MR project) and a less-risky project (LR project). Subjects were informed about the mean ( $\mu_{LR\ project} = 7.9\%$ ;  $\mu_{MR\ project} = 9.7\%$ ) and standard deviation of the project returns ( $\sigma_{LR\ project} = 2.5\%$ ;  $\sigma_{MR\ project} = 4\%$ ) and were instructed that the MR project was riskier than the LR project since it was associated with a higher standard deviation<sup>4</sup>. Subjects were also told that the returns from both projects were uncertain and that the actual returns would only become known after they made their allocation decision. All subjects started with an initial endowment of \$10. Consistent with the model setup, the study comprised of three periods: a baseline period (8 rounds), a treatment period (24 rounds) and a post-treatment period (8 rounds). The complete set of task instructions are provided in Appendix A2.

In every round, subjects observed a graph showing the performance of each project during the past 10 rounds and then allocated their entire current endowment between the two projects based on their forecasted returns for the upcoming period. An example of the stimuli used in the study is shown in Appendix A3. 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 shown the same forecast

---

<sup>4</sup> When we compare the coefficient of variation (i.e., ratio of standard deviation to the mean), we see that the MR project is about 1.3 times riskier than the LR project.

of the returns for the upcoming period, but the suggested allocation differed based on the algorithm's level of risk. Figure 3 below shows the experimental design for the study.

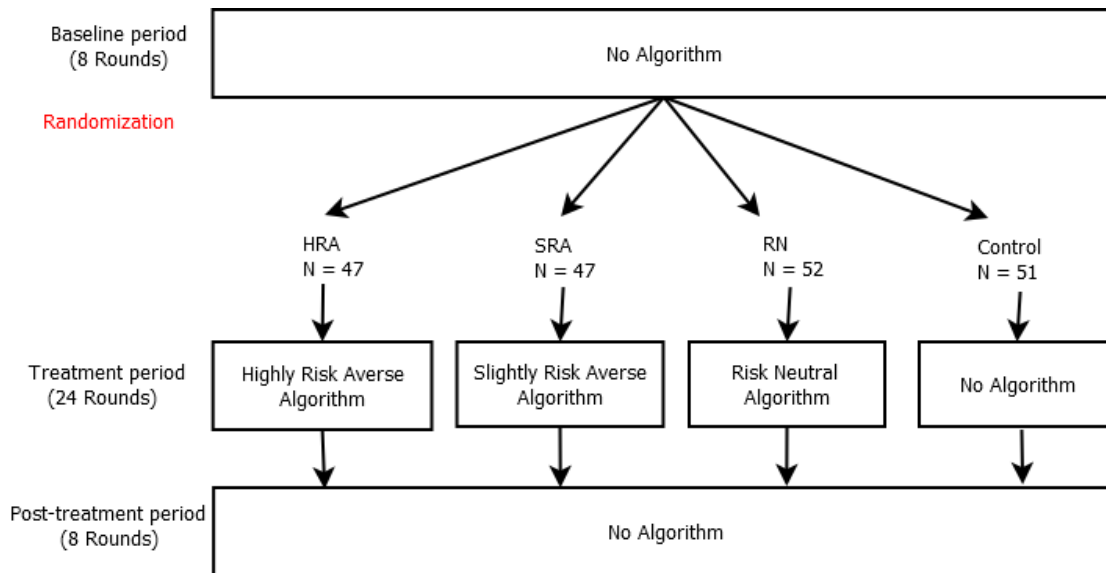


Figure 3. Study design

I informed subjects in the treatment conditions (HRA, SRA and RN) that the algorithmic recommendation was the output of a mathematical model. Importantly, they were briefed about the level of risk associated with the algorithmic recommendation. Specifically, subjects in the RN condition learned that the algorithm would be *risk neutral*, in that the algorithmic recommendation maximized expected returns without any consideration of risk. Subjects in the SRA (HRA) condition learned that the algorithm would be *slightly (highly) risk-averse*, meaning that the algorithm recommended allocating more resources to the less-risky project compared to a risk neutral allocation<sup>5</sup>. Subjects in all the treatment conditions were free whether or not to

---

<sup>5</sup> The distinction between slightly risk-averse (SRA) and highly risk-averse (HRA) algorithm was made by using different values of risk aversion parameter ( $\lambda$ ). We used a  $\lambda$  value of 0.5 for designing the

follow the algorithmic recommendation. Details regarding the algorithmic design are provided in Appendix A4.

At the end of each round, subjects received extensive feedback about their performance. Specifically, they viewed their cumulative wealth as well as the returns they earned along with two counterfactuals: the returns they could have earned had they followed the algorithmic recommendation<sup>6</sup> and the maximum possible returns for that round. An example of the feedback screen employed in the study is shown in Appendix A3. At the end of the treatment period, I removed the algorithmic aid for all subjects, and they proceeded to complete the final 8 rounds of the study. 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.

### 3.2 Results

**Resource Allocation Decisions.** In the study, I use the proportion of the endowment (i.e., resources) allocated to the MR project in a given round as the key dependent variable measure. I examine resource allocation decisions in each of the three experimental periods separately. In Figure 4, I observe no discernible differences in resource allocation decisions among subjects across different conditions during the baseline period, due to random assignment (also, see Table 1, columns 1-2). Furthermore, consistent with hypothesis H1, I observe that subjects in the HRA condition acted more risk-averse than those in the control condition during the treatment as well as the post-treatment periods. For example, subjects in the

---

SRA algorithm in both chapters 1 and 2. We used  $\lambda$  value of 4 and 2 for designing the HRA algorithm in chapter 1 and 2, respectively.

<sup>6</sup> Algorithmic returns were displayed only for subjects assigned to either HRA, SRA or RN conditions during the treatment period.

HRA condition allocated 9.3% and 7.4% less resources (on average per round) to the MR project compared to those in the control condition during the treatment and post-treatment period, respectively (see Table 1, columns 3-6). The results in the post-treatment period were robust when controlling for the returns earned during the treatment period (see Appendix A5). Thus, subjects in the HRA condition changed their resource allocation decisions in response to encountering a highly risk-averse algorithm and these behavioral changes were sticky even after the algorithm was removed. Hence, the study offered initial empirical evidence in support of hypotheses H1 and H2. Moreover, I found that subjects in the SRA condition behaved similarly to those in the control condition. While subjects in the RN condition were more risk seeking than those in the control condition during the treatment period, they did not significantly differ from those in the control condition during the post-treatment period.<sup>7</sup>

---

<sup>7</sup> Thanks to Prof. Kriti Jain for identifying that the control group is not stable through the course of the study (i.e., there is a small upward trend towards more allocating more resources to the MR project). This issue is rectified in chapter 2 and the results remain consistent.

Note: Shaded region signifies  $\pm 1$  standard error




Figure 4. Endowment Allocated to MR Project

**Table 1. OLS regression on resource allocation decisions with round-fixed effects included**

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

	<i>Dependent variable:</i>					
	Endowment Allocated to MR Project					
	Baseline		Treatment		Post-Treatment	
	(1)	(2)	(3)	(4)	(5)	(6)
Highly Risk-averse	0.536 (2.951)	-0.180 (3.070)	-9.232*** (2.482)	-9.271*** (2.076)	-8.787** (3.617)	-7.443** (3.308)
Risk Neutral	0.292 (2.824)	0.037 (2.992)	11.916*** (2.609)	11.002*** (2.380)	-2.277 (3.367)	-1.871 (3.455)
Slightly Risk-averse	2.257 (3.060)	2.231 (3.272)	-1.226 (2.398)	-2.590 (2.026)	-4.456 (3.237)	-4.402 (3.177)
Age		0.050 (0.354)		-0.146 (0.229)		-0.401 (0.374)
Gender		1.988 (2.201)		-0.218 (1.579)		-0.167 (2.278)
Previous Round Returns		-4.230*** (1.391)		3.014*** (0.726)		4.338** (2.075)
Baseline Allocation				0.329*** (0.054)		0.527*** (0.091)
Constant	57.729*** (2.394)	87.114*** (15.189)	59.045*** (2.492)	28.876*** (7.804)	75.316*** (2.680)	5.384 (24.745)
Observations	1,576	1,372	4,728	4,508	1,576	1,372
Adjusted R <sup>2</sup>	0.034	0.287	0.103	0.128	0.030	0.130

*Note:* \*p<0.1; \*\*p<0.05; \*\*\*p<0.01

**Advice Utilization.** Subjects in the HRA condition showed the largest deviation from algorithmic recommendation as they allocated on average 22% more resources to the MR project than what was suggested by the HRA algorithm. Subjects in the SRA condition were also significantly less risk-averse than the (SRA) algorithmic recommendation, albeit by a smaller magnitude (see Appendix A6).

**Earnings.** The returns earned by subjects in all conditions during each of the three experimental periods is reported in Appendix A7. Subjects in RN and SRA

conditions earned slightly more per round ( $\sim 0.2\%$ ) than those in the HRA condition during the treatment period ( $t(95.46) = 4.48, p < 0.001$  and  $t(91.33) = 4.16, p < 0.001$ , respectively). However, I did not find any significant difference in returns earned by subjects across conditions during the post-treatment period. Additionally, there were no significant differences in the standard deviation of returns earned by subjects across conditions during treatment or post-treatment periods.

#### **4. Discussion**

Using the literature on the anchoring-and-adjustment heuristic, I predicted that DMs' resource allocation decisions change when exposed to an algorithmic aid and the changes stick even after the algorithmic aid is removed. I tested the empirical validity of these claims in a randomized experiment using a portfolio management task. In the study, I showed that DMs in portfolio optimization task exhibited higher levels of risk aversion in their investment choices subsequent to observing recommendations from a highly risk-averse algorithm. Moreover, DMs continued to exhibit higher levels of risk aversion even after the algorithmic recommendation was removed.

The study has direct implications for the current literature on the anchoring and adjustment heuristic (Tversky & Kahneman, 1974). Unlike recent studies that have demonstrated the persistence of anchoring using "uninformative" anchors (i.e., random values which are uncorrelated with the "true value" of the estimate), I show that highly risk-averse algorithms, which contain informational value, can leave an imprint on the DM following temporary exposure (Yoon & Fong, 2019). This finding is particularly significant given that DMs are aware of the level of risk aversion of the algorithmic recommendation they receive and have complete freedom in their

decision-making (i.e., they are free to choose whether or not to follow to algorithmic recommendation).

Furthermore, I contribute to the line of research showing that within-task risk attitudes can change. Existing studies that incorporate standard models of risk aversion (Pratt 1978) commonly employ a static parameter for modelling individual risk attitudes. However, a few prior studies have shown that subjects' risk and ambiguity attitudes may change not only across decision tasks (Fox & Tversky, 1995); (Abdellaoui et al., 2011) but also within the same task due to learning or as a result of varying dimensions such as the size or type of incentives (Blavatsky et al., 2021). Using a repeated task design, I demonstrate that within-task risk attitudes can change in response to algorithmic recommendations.

I believe that the study has important domain-specific implications. For example, in the field of finance, subscription-based decision support is becoming increasingly common not only for professional traders, but also for lay people investing their own private wealth. The results show that in these contexts, companies could use algorithmic recommendations to nudge employees or consumers towards a specific risk-taking behavior. In fact, governments are trying to increase consumer participation in equity markets (Mehra & Prescott, 1985). In such contexts, temporary algorithmic support could be used to reduce risk aversion and increase investments in equities.

I next discuss some of the study limitations. First, I designed the experimental study in order to investigate risky decisions in the domain of gains. In particular, I focused on MR and LR projects with positive mean expected return and small standard deviations. Therefore, losses in the settings did not occur. However, it is possible that the findings do not extend to resource allocation contexts involving losses. Future

research could investigate how people respond to algorithmic recommendations when exposed to losses.

Second, I restricted myself to studying algorithms that assumed either risk neutrality or risk aversion. I did not explore risk seeking algorithms, as they are hard to explain to subjects. Given that a risk neutral recommendation maximizes expected profits, communicating why a risk seeking strategy could potentially result in even higher profits is not straightforward. However, the effect of risk seeking algorithms which, for instance, maximize some arbitrary upper percentile of the outcome distribution, should be tested in future research.

Finally, while the results were demonstrated in a portfolio management task in the domain of finance, it is unclear whether they will hold in other resource allocation problems such as inventory management in operations. If the results are in fact domain-invariant, it would help generalize the findings regarding risk-averse algorithmic influence across a broader spectrum of resource allocation problems. Moreover, a successful replication would bolster the robustness of the original findings and aligns with the recent trend within social-sciences to increase the transparency and reliability of experimental research (Davis et al., 2023). I attempt this exercise of self-replication in chapter 2.

## Chapter 2: Risk-averse algorithmic support in inventory management

**Abstract:** Chapter 2 serves as a conceptual replication of the findings in chapter 1 in the domain of inventory management. I examine how order decisions change when decision-makers (DMs) are temporarily exposed to algorithmic aids that are programmed with predetermined levels of risk aversion in a multi-item newsvendor task. I find that participants exposed to highly risk-averse (HRA) algorithmic recommendations tend to make more risk-averse order decisions, even though they deviate from the advice. Strikingly, the changes in order decisions persist even after the (HRA) algorithm is removed. Thus, chapter 2 further demonstrates the mutability of human behavior when temporarily exposed to risk-averse algorithmic aids.

**Keywords:** algorithm | risk aversion | anchoring | newsvendor

## 1. Introduction

The newsvendor model remains a foundational topic of study in inventory management. The enduring academic interest in the newsvendor model may be attributed to the problem's simplicity, the model's flexibility, and its practical relevance across diverse domains. Schweitzer and Cachon (2000) invigorated behavioral research in the newsvendor problem by documenting the "pull-to-center" (PtC) bias, which is the tendency of decision-makers (DMs) to place orders that lie between mean demand and optimal order quantity. While there is considerable heterogeneity in ordering behavior (Lau et al., 2014), most DMs do not order the risk-neutral expected profit maximizing quantity. The existing literature commonly treats such a deviation from the normative benchmark as a judgmental bias and has frequently attempted to improve DMs' order behavior through the provision of risk-neutral, profit maximizing decision support. For instance, Lee and Siemsen (2017) provide risk-neutral order recommendations and find that it reduces the PtC bias. This finding is echoed in similar studies, which suggest that risk-neutral order recommendations are likely to be a useful means for debiasing decision behavior (Feng & Gao, 2020; Zhang & Siemsen, 2019).

However, such a normative order strategy may be misaligned with the DM's actual preference for taking risks which may result from both individual differences in risk attitudes as well as an institutional requirement to pursue conservative order strategies. 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 is expected that DMs would be cautious or risk-averse to avoid chances of a loss or

unfavorable profit levels (Chen et al., 2007b). In such cases, DMs might prefer risk-averse decision support. In fact, in similar resource allocations tasks such as the portfolio optimization in finance, algorithms calibrated with a particular level of risk aversion are commonly used (D'Acunto et al., 2019).

What does it mean to be risk-averse in the newsvendor context? In the classic newsvendor model, Eeckhoudt et al. (1995) showed that a hypothetical risk-averse newsvendor would order lower quantities than their risk-neutral counterpart and experience profits that are less sensitive to variations in demand. Since then, there has been considerable analytical work that incorporates risk aversion in inventory models using different objective functions (e.g., Ahmed et al., 2007; Prakash Katariya et al., 2014), varying the assumptions regarding the newsvendor cost structure (e.g., Wang et al., 2009; Wu et al., 2009), exploring different time horizons (e.g., Chen et al., 2007) and extending the analysis from single to multiple products (e.g., Choi & Ruszczyński, 2011). However, surprisingly, commensurate behavioral work on risk aversion has been extremely limited. To my knowledge, only a few papers focus explicitly on risk aversion in newsvendor experiments (Becker-Peth et al., 2018; De Véricourt et al., 2013). 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, I add to the literature on risk aversion in inventory management by studying how people with different risk attitudes respond to algorithms of different risk aversion levels.

Furthermore, unlike previous studies on newsvendor decision support, I study *algorithmic withdrawal* i.e., how an algorithmic decision support impacts newsvendor decisions even after they are no longer available. I believe this is a critical issue to examine for several reasons, as already discussed in Chapter 1. Again, relying on the anchoring and adjustment literature, I derive and test the same predictions regarding changes in resource allocation decisions in response to algorithmic recommendations in a multi-item newsvendor task. I specifically choose the constrained multi-item newsvendor context to replicate the results in Chapter 1, since the task is cognitively challenging, and representative of real-world inventory problems. Also, the necessity to allocate capacity (i.e., shelf-space) among multiple alternatives (i.e., products with different demand volatilities) makes the concept of risk highly salient from the perspective of the DM in the multi-item newsvendor task.

Additionally, I made five important experimental design changes to address potential limitations from the previous study in chapter 1. First, in chapter 1, the RN algorithm allocated 100% of the endowment to the MR project in many rounds, which made it impossible to be more risk seeking than the RN algorithm. I corrected this design issue. Second, I chose neutral labels of product A and product B to denote the less-risky (LR) and more-risky (MR) alternatives respectively in the multi-item newsvendor study. Third, I extended the post-treatment period to 12 rounds (vs. 8 rounds earlier) in order to test whether changes in ordering behavior would persist for a longer duration after the algorithmic recommendation is removed. Fourth, I ensured that demand distributions of products A and B were identical in the baseline period and in the post-treatment period to allow for direct within-subject comparison of

treatment effects<sup>8</sup>. Finally, I provided better individual performance feedback using a combination of a table and a graph as opposed to only using a table in chapter 1.

In my experimental study, I again 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 and hedging away from it. However, the modified (risk-averse) order decisions helped subjects achieve more stable profits.

## 2. Multi-item newsvendor model

**Setup.** A DM in my inventory management 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  $c$ ) and are *high-profit* 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  $\mu_{At}$  and  $\mu_{Bt}$  denote the mean demand for products A and B and  $\sigma_{At}$  and  $\sigma_{Bt}$  represent the standard deviation of demand for products A and B for time period  $t$ . In my 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.,  $\frac{\sigma_{Bt}}{\mu_{Bt}} = 2 * \frac{\sigma_{At}}{\mu_{At}} \forall t$ .

---

<sup>8</sup> Since there are unequal number of rounds in the baseline and post-treatment period in the study, I made the demand distributions in 8 rounds of the baseline period identical to the demand distributions in 8 rounds of the post-treatment period.

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 my multi-item newsvendor setup.

Random profits for each product (denoted by  $\pi_{At}, \pi_{Bt}$ ) in time period  $t$  are calculated using Eqs. (3) and (4) below. The total random profit ( $\pi_t$ ) is the sum of the individual profits of product A and B.

$$\pi_{At} = p \min(q_{At}, d_{At}) - cq_{At} \quad (3)$$

$$\pi_{Bt} = p \min(q_{Bt}, d_{Bt}) - cq_{Bt} \quad (4)$$

$$\pi_t = \pi_{At} + \pi_{Bt}. \quad (5)$$

**Algorithm design.** I design the algorithm such that it tries to maximize an objective function  $F$ , which can be expressed as follows:<sup>9</sup>

$$F = E(\pi_t) - \lambda * V(\pi_t) \quad (6)$$

In Eq. (4),  $E(\pi_t)$  is the expected profits,  $V(\pi_t)$  represents the variance in profits and  $\lambda$  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  $\lambda$  should lead to less ordering of*

---

<sup>9</sup> The objective function of the form  $F = E - \lambda V$  can be derived by assuming a DM tries to maximize an exponential utility function with a particular level of risk aversion  $\lambda$ . While newsvendor profits are not normally distributed, this formulation of  $F$  suffices to generate intuitively explainable risk-averse algorithmic recommendations (For an illustration, refer to Appendix B.1).

the riskier product B. For the study parameters I used, I simulated this result, and it is shown in Appendix B1.

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 the experimental study, I chose  $\lambda$  values of 0.5 and 2 to model the SRA and HRA algorithms, respectively. Note that  $\lambda = 0$  for a RN algorithm.<sup>10</sup> I confirmed that  $q_{RN}^* > q_{SRA}^* > q_{HRA}^*$  in the study for all decision rounds.

**Predictions.** See Chapter 1 on how I derive the predictions using anchoring paradigm. The same hypotheses are tested in Chapter 2. As noted earlier, such an exercise in self-replication in a different context with a different resource will broaden the applicability of my findings and bolster the reliability of the original results.

### 3. Experimental Study

#### 3.1 Methods

**Subjects.** Subjects were Master's degree students from a large private university. The study was conducted online, and subjects earned class credits in return for their participation. In addition to earning extra course credits, I selected 10 subjects at random and awarded them Amazon vouchers matching the profits of one randomly chosen round.<sup>11</sup> The average earnings were \$22 per winning subject. The final sample size was 184 (female = 40%) with an average age of 27 years.

**Design and Procedure.** Once subjects consented to take part in the study and were briefed of the incentives, they were asked to assume the role of a manager of a

---

<sup>10</sup>  $q_{RN}^*$  can also be obtained in every time period from the heuristic proposed by Erlebacher (2000), where  $q_{RN}^* = \mu_{Bt} + \frac{\sigma_{Bt}}{\sigma_{At} + \sigma_{Bt}}(100 - \mu_{At} - \mu_{Bt})$ .

<sup>11</sup> This variation of the random incentive system is almost the exclusively used incentive system in similar individual choice experiments involving risk preferences (See van de Kuilen & Wakker, 2011).

retail store which sold products A and B. Subjects were informed about the multi-item news vendor set-up described in Section 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 is provided in Appendix B2. The demand distributions for products A and B are listed in Table 2 below.

**Table 2. Demand distribution of products A and B**

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

In Table 2,  $z_t$  is a random integer between 0 and 12, independently generated for every round  $t$  (and undisclosed to subjects). I 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 ( $\sigma = 10.1$ ) was lower than the demand volatility of product B ( $\sigma = 36.1$ ). 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 the 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. Figure 5 below summarizes the experimental design for the study.

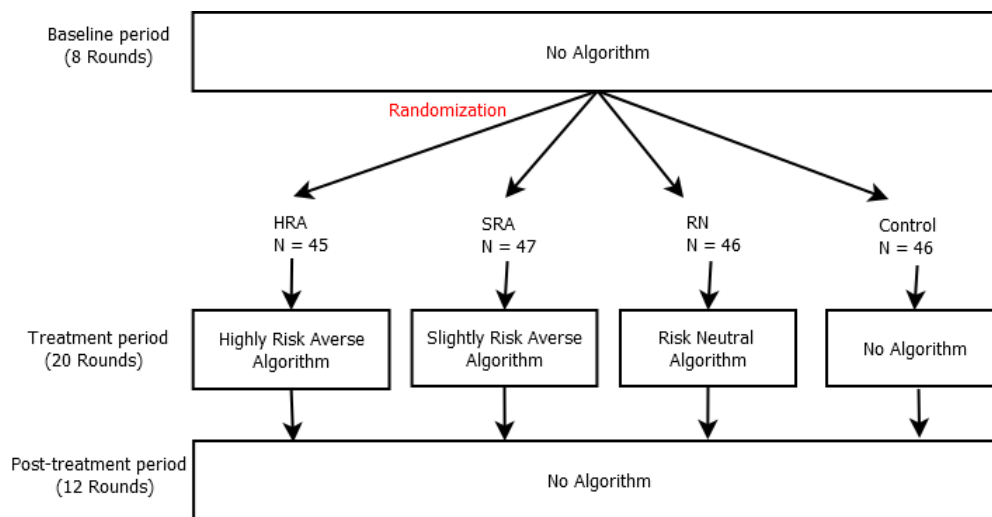


Figure 5. Study Design

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.<sup>12</sup> In the graph, subjects viewed their actual profits along with two counterfactuals: the profits they would have received had they followed the

<sup>12</sup> The algorithmic recommended order quantity was only displayed to subjects in the RA algorithm condition during the treatment period.

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 subjects is shown in Appendix B3. At the end of the treatment period, I removed the algorithmic aid for all subjects, who then proceeded to complete the final 12 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.

### **3.2 Results**

**Ordering behavior.** I operationalized ordering behavior as the number of units of (riskier) product B ordered per round. I analyzed ordering behavior in each of the three experimental periods separately. In Figure 6, we observe no discernible differences in ordering behavior among subjects across different conditions during the baseline period due to random assignment (also see Table 3, columns 1-2). 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 3, columns 3-4). For example, subjects in HRA, SRA and RN conditions ordered, on average, 7.1, 4.7 and 3.8 fewer units of product B, respectively, compared to subjects in the control condition during the treatment period. In Figure 6, 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 3, 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, 4.4 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 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, the study provides support for H1 and H2.

*Note: Shaded region signifies  $\pm 1$  standard error*

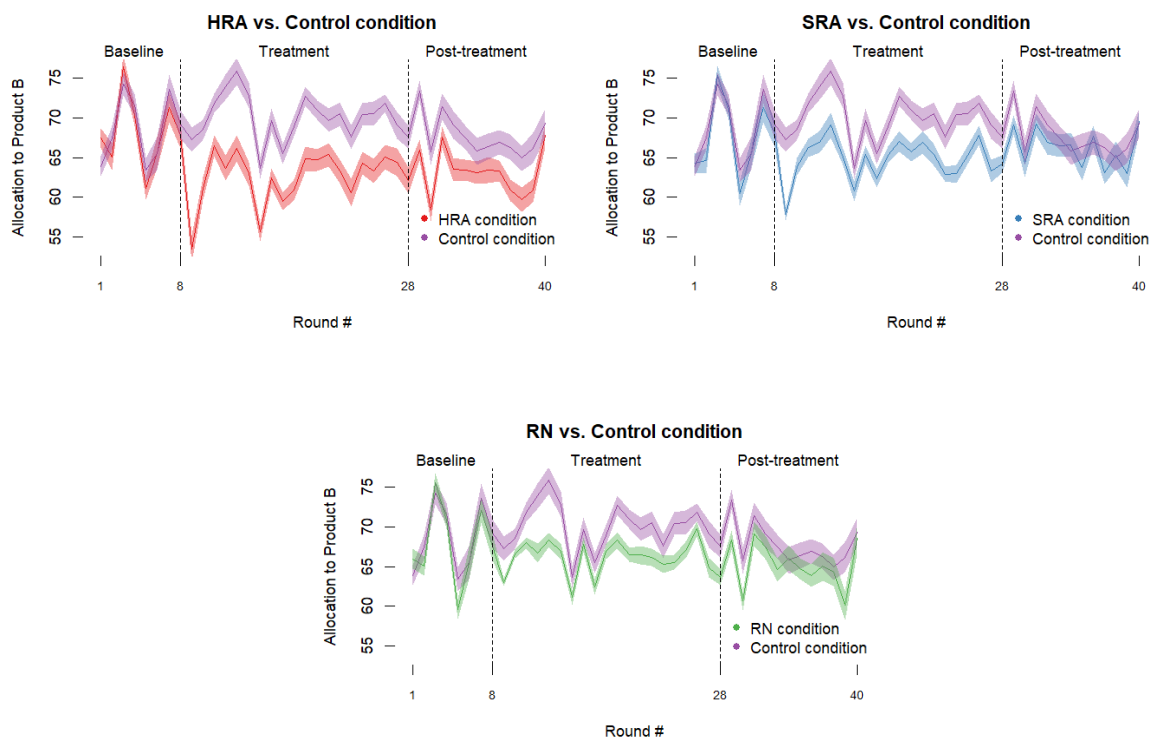


Figure 6. Allocation to Product B

**Table 3. OLS regression on ordering behavior with round-fixed effects included**

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

	<i>Dependent variable:</i>					
	Allocation to Product B					
	Baseline		Treatment		Post-Treatment	
	(1)	(2)	(3)	(4)	(5)	(6)
Highly Risk-averse	-0.486 (1.259)	-1.101 (1.297)	-7.415*** (1.074)	-7.107*** (0.948)	-4.582*** (1.201)	-4.373*** (1.182)
Risk-neutral	-0.852 (1.202)	-1.334 (1.223)	-3.870*** (0.833)	-3.828*** (0.737)	-2.513** (1.186)	-2.212* (1.135)
Slightly Risk-averse	-1.042 (1.246)	-1.183 (1.302)	-5.216*** (0.933)	-4.727*** (0.818)	-1.496 (1.367)	-0.916 (1.332)
Age		-0.007 (0.125)		-0.007 (0.089)		0.050 (0.118)
Gender		0.179 (0.892)		1.153* (0.628)		1.345 (0.928)
Previous Round Profits		0.014*** (0.005)		0.002 (0.003)		0.002 (0.004)
Baseline Allocation				0.348*** (0.069)		0.388*** (0.076)
Constant	65.992*** (1.046)	55.261*** (5.616)	64.559*** (0.988)	43.488*** (6.062)	71.327*** (1.050)	35.661*** (6.364)
Observations	1,472	1,288	3,680	3,496	2,208	2,024
Adjusted R <sup>2</sup>	0.139	0.388	0.146	0.173	0.066	0.097

Note: \*p<0.1; \*\*p<0.05; \*\*\*p<0.01

**Advice Utilization.** Subjects in the HRA condition showed the largest deviation from algorithmic recommendation as they 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 expectedly smaller (See Appendix B4). Consistent with previous studies, I also computed a weight on advice (WOA) score, which is a measure commonly used to capture the extent of advice utilization (Logg et al., 2019; Soll & Larrick, 2009). I discuss the WOA measure along with study 2a results in Chapter 3 to

facilitate a meaningful comparison of advice utilization between endogenous versus exogenous algorithmic choice.

**Average profits.** The average profits earned by subjects in all conditions during each of the three experimental periods are reported in Appendix B4. 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, I did not find any significant difference in average profits earned by subjects across different conditions during the post-treatment period.

**Variability in profits.** I measured variability in profits as the average standard deviation of profits, computed for each subject across treatment and post-treatment periods, separately. From Figure 7 below, I 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, Figure 7 below).

*Note: Error bars signify  $\pm 1$  standard error*

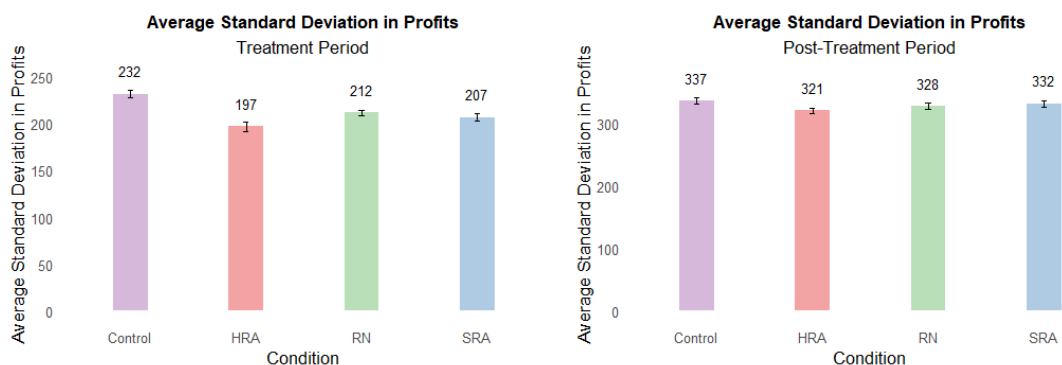


Figure 7. Average Standard Deviation in Profits

Overall, I find that changes in order decisions are strong when the task is cognitively complex and when the algorithmic anchor is extreme (i.e., highly risk-averse). I also find persistence of anchoring effects to be strong for subjects who faced the HRA algorithm, and I explore the underlying reasons in sub-section 3.3.<sup>13</sup>

### **3.3 Persistence of treatment effects.**

I examine the reasons for change in order decisions in the post-treatment period compared to the baseline period. Before I test the anchoring paradigm, I consider a few alternate explanations, such as 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 Appendix B5). Nonetheless, if anchoring effects persist, I 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). 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 my 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.

In order to test this proposition, I regressed the “change in the average order of product B” during the post-treatment period (vis-à-vis baseline)  $(\bar{q}_2 - \bar{q}_0)$ ” onto the “change in the average order of product B during the treatment period (vis-à-vis

---

<sup>13</sup> We do not examine behavioral biases bias since order decisions for each product are not independent in the multi-item newsvendor task.

baseline)  $(\bar{q}_1 - \bar{q}_0)$ ” for subjects in both HRA Algorithm and control conditions in study 1. However, since choices during the treatment period might be correlated with unobservables that predict choices during the post-treatment period, simple OLS regression is not suitable. Therefore, I run an instrumental variable (IV) regression using random assignment to the experimental condition as the instrument. I believe that condition is a useful instrument because it predicts choices during the treatment period (relevance) while being uncorrelated with unobservables that predict choices during the post-treatment period (exclusion restriction). The relevance criterion for this instrument was empirically verified with an F-test (Test-statistic = 50.1;  $p < 0.001$ ) confirming that condition is a strong instrument. The results of the IV regression are reported in Table 4. I find that subjects who ordered fewer units of product B during the treatment period (vis-a-vis baseline) also ordered fewer units of product B during the post-treatment period (vis-a-vis baseline).

**Table 4. IV regression on mean change in allocation to product B**

	<i>Dependent variable:</i>
	$(\bar{q}_2 - \bar{q}_0)$
$(q_1 - \bar{q}_0)^{predicted}$	0.608*** (0.148)
Age	0.049 (0.159)
Gender	0.114 (1.154)
Baseline Allocation	-5.730* (3.275)
Constant	1.284 (4.762)
Observations	91
Adjusted R <sup>2</sup>	0.452

*Note:* \*p<0.1; \*\*p<0.05; \*\*\*p<0.01

## 4. Discussion

I study individual decision-making in resource allocation tasks when temporary algorithmic support of varying levels of risk is available. In the experimental study, I 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, I showed that DMs facing highly risk-averse algorithms who ordered less of the riskier product during the treatment period were the ones who continued to do so in the post-treatment period.

A key observation from the study 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. Thus, the least used advice resulted in the most significant changes in behavior. How do we explain this paradox? First, from an experimental design perspective, the low levels of advice utilization for HRA algorithm can be explained from the fact that subjects were fully aware of the algorithm's level of risk aversion and had access to extensive 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 (or anchors) that are farthest away from their own estimates, a phenomenon more commonly referred to as egocentric discounting (Rader et al., 2017). Thus, given that HRA algorithmic recommendation represents the most extreme type of advice (or anchor), it is understandable that subjects deviated a lot from the HRA recommendation (in a *less* risk-averse direction as a way to hedge). However, the shift was not enough (i.e., the insufficient adjustment in anchoring-and-adjustment

heuristic) to return to their “baseline” ordering behavior. This is the reason we observe significant changes in ordering behavior for subjects who observed the HRA algorithm, despite low advice adherence. Additionally, I explain the persistence of change in ordering behavior for DMs who observe HRA algorithms using theories on habit formation (Becker & Murphy, 1988). In my 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, I 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).

This 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. Specifically, while I observed significant changes in ordering behavior with a highly risk-averse algorithm, I do not recommend consistently increasing the level of risk aversion since the results suggest that trust and advice utilization decrease as the level of risk aversion increases.

While chapter 2 further helps demonstrate the influence of highly risk-averse algorithms in resource allocation problems, it is useful to reflect on two pertinent questions. First, what are the boundary conditions of the observed risk-averse algorithmic effects? Second, can risk-averse algorithms be used to *objectively* improve decision-making in inventory management? I address both these critical questions in chapter 3.

## Chapter 3: Risk-averse algorithmic support in inventory management: Further Extensions

**Abstract:** I explore whether critical factors such as decision autonomy (i.e., choosing an algorithm of a particular risk level versus random assignment) and source of advice (human versus algorithm) affect the extent of algorithmic influence in inventory management decisions. I find that algorithmic advice utilization is similar regardless of whether the algorithm is externally assigned or chosen by the subjects themselves. Furthermore, when I compare algorithmic recommendations with equivalent human advice, I observe that subjects do not exhibit any algorithm aversion. Finally, I also demonstrate that risk-averse algorithmic support can help counteract decision biases in the low-profit margin regime in the single-item newsvendor problem. Consequently, organizations can deploy risk-averse algorithmic tools to improve inventory decisions while preserving managerial autonomy.

**Keywords:** algorithm aversion | newsvendor | risk aversion | pull to center bias

## 1. Introduction

Recent research on algorithm aversion and advice-taking suggests that there could be differences in algorithmic utilization depending on whether the algorithmic assignment is endogenous or exogenous (Dietvorst et al., 2018; Rader et al., 2017) and whether the recommendations are from an algorithm or a human (Dietvorst et al., 2015). Therefore, the primary goal of chapter 3 is to examine whether critical factors, such as decision autonomy (i.e., choosing an algorithm versus random assignment) and source of advice (human versus algorithm) have an impact on the extent of algorithmic influence. In fact, in a recent review of 30 years of behavioral intervention studies in mainstream operations journals, the authors point to the dearth of studies that assume a voluntarist behavior (i.e., individuals having agency instead of their actions being determined exogenously)(Franco et al., 2021). Thus, by explicitly providing decision autonomy (i.e., agency to choose an algorithm), I hope to fill this gap in the current literature on behavioral OR studies. Importantly, examining factors regarding *autonomy* and *source* helps shed light on the boundary conditions of the observed algorithmic influence in inventory management tasks.

A second important agenda of chapter 3 is to explore the behavioral implications of using risk-averse algorithmic aids in the classic, single-item newsvendor task. In particular, I want to examine if risk-averse algorithms can be used to mitigate the Pull-to-Center (PtC) bias, which is the tendency of decision-makers (DMs) to place orders that lie between mean demand and optimal order quantity (Schweitzer & Cachon, 2000). This investigation is important since previous studies on algorithmic support do not examine if algorithms can influence order decisions even *after the algorithm is no longer available*. (Feng & Gao, 2020; Lee & Siemsen, 2017; Zhang & Siemsen, 2019).

Again, relying on the anchoring literature, I first derive predictions regarding how order decisions change in response to algorithmic recommendations and test the predictions across two experimental studies involving a single- and multi-item newsvendor task (Studies 1, 2a and 2b, respectively). In study 1, I used a risk averse algorithmic recommendation in the classic single-item newsvendor task. In studies 2a and 2b, I employed multiple algorithms with varying levels of risk aversion and investigated the boundary conditions of the observed anchoring effects regarding decision autonomy (i.e., choosing an algorithm of a particular risk level versus random assignment) and source of advice (human versus algorithm) in the multi-item newsvendor task. All studies were comprised of three experimental periods: a baseline period, a treatment period, and a post-treatment period and the experimental protocol is similar to those employed in chapters 1 and 2.

Across all studies, I found that subjects who received risk averse algorithmic recommendations ordered less of the focal (i.e., riskier) product and continued to do so even after the algorithmic recommendation was removed. The modified order decisions helped subjects achieve more stable profits and also minimized the PtC bias in the single-item newsvendor task (study 1). Additionally, in the two follow-up studies, I found that changes in ordering behavior and anchoring effects were similar regardless of whether the algorithm was externally assigned (i.e., exogenous) or chosen by subjects themselves (i.e., endogenous) (study 2a), and subjects did not exhibit any algorithm aversion when I compared their responses to highly risk-averse algorithmic recommendations and equivalent human advice (study 2b).<sup>14</sup>

---

<sup>14</sup> Additionally, in study 2a, we found that over a third of the subjects preferred the HRA algorithmic support when they had the autonomy to choose between a HRA and a RN algorithmic support.

## 2. Single-item newsvendor model

**Setup.** A DM in the classic single-item newsvendor problem sells a product with known cost structure (i.e., unit selling price of  $p$  and unit cost price of  $c$ , where  $p > c$ ). I consider a *low-profit* margin product with critical fractile  $< 0.5$  (Schweitzer & Cachon, 2000a).<sup>15</sup> Low-profit margin products have also been used extensively in previous newsvendor experiments (see Zhang and Siemsen (2019), Table 2 for a comprehensive list of studies). The demand for the product is uncertain and varies across different time periods. The actual demand values are drawn from a discrete uniform distribution  $\sim U(0, D_t)$ , which is known to the DM at the start of time period  $t$ . Let  $d_t$  denote the actual demand in a particular time period  $t$ . Additionally, the DM is aware that actual demand values are independent across different time periods. Let  $\mu_t$  denote the mean demand and  $\sigma_t$  represent the standard deviation of demand. The DM has to decide the order quantity (denoted by  $q_t$ ) in every time period. Importantly, actual demand values will be realized only *after* the order decision is made. There are no salvage costs (from resale of excess unsold units) or stockout costs (penalty for running out of stock) included in the newsvendor model setup. The random profits (denoted by  $\pi_t$ ) at the end of time period  $t$  are calculated using Eq. (7) below.

$$\pi_t = \{p \cdot \min(q_t, d_t) - cq_t\} \quad (7)$$

**Algorithm design:** The risk-averse algorithm is designed to maximize expected utility, subject to a degree of risk aversion  $\lambda$ . I assume an exponential form of the utility function  $U_t$ , which is common in the newsvendor literature (Eeckhoudt et al., 1995).

---

<sup>15</sup> Critical fractile is calculated using the formula  $(p - c)/p$ .

$$U_t = \begin{cases} -\frac{e^{-\lambda\pi_t}}{\lambda}, & \lambda \neq 0 \\ \pi_t, & \lambda = 0 \end{cases} \quad (8)$$

The utility function  $U_t$  can accommodate both risk-averse ( $\lambda > 0$ ) and risk seeking ( $\lambda < 0$ ) behaviors. Higher positive values of  $\lambda$  indicate more risk version. For a risk-neutral algorithm or DM ( $\lambda = 0$ ), the optimal order quantity ( $q_t^*$ ) is the well-known expression  $q_t^* = \frac{(p-c)}{p}D_t$ . It can be further seen that optimal order quantity for a risk-averse algorithm is a strictly decreasing function of  $\lambda$  (For proof, see Eeckhoudt et al., (1995). For further details on the algorithm design, see Appendix C1).

**Predictions.** The objective is to examine how DMs modify their order decisions in response to risk-averse algorithmic recommendations and derive predictions for the experiment. Without loss of generality, I again 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 a *RA Algorithm* condition make order decisions after observing a risk-averse algorithmic recommendation of an optimal order 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. I allow  $q$  in the 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. I assume that  $q_t$  evolves as follows:

**Baseline period:** As subjects are randomly assigned, I do not expect to see any difference in (average)  $q_0$  across conditions at the end of the baseline period i.e.,

$q_0^{RA\ Algorithm} = q_0^{Control} = q_0$ . As observed in previous newsvendor studies, I expect subjects in the baseline period to order in excess of  $q_0^*$  (on average) in the *low-profit* regime (Zhang & Siemsen, 2019).

**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 that she encounters. In the presence of an algorithm, I assume that the DM's final order  $\bar{q}_1$  is a weighted average of her baseline order  $q_0$  and the risk-averse algorithmic recommendation  $q_{RA}$  ( $q_{RA} < q_0^* < q_0$ ). Therefore, using the anchoring and adjustment heuristic, the DM's final order during the treatment period is given by:

$$\bar{q}_1 = \beta q_0 + (1 - \beta)q_{RA} \quad (9)$$

where  $0 \leq \beta \leq 1$ .

In Eq. (9),  $\beta$  captures the weight DMs give to baseline order  $q_0$  and the adjustment they make away from the anchor  $q_{RA}$ . It follows from Eq. (9) that DMs who observe the RA algorithm would order less than those in the control condition (since  $\bar{q}_1 < q_0$ ).

**H1:** *DMs who observe a risk-averse algorithm order less compared to those who do not observe any algorithm.*

**H2:** *DMs who observe a risk-averse algorithm continue to order less even after the algorithm is removed.*<sup>16</sup>

**Pull to Center bias.** Several previous newsvendor studies have documented that DMs tend to place orders between mean demand and optimal levels. This is

---

<sup>16</sup> H2 is based on the rationale of persistence of anchoring (See analogous section in Chapter 1)

called pull-to-center (PtC) bias and can be quantified as follows for every subject (Zhang & Siemsen, 2019)

$$PtC\ score = \frac{q - q^*}{\mu - q^*} \quad (10)$$

PtC score is a fraction typically between 0 and 1 with higher values indicating more bias (i.e., orders close to mean demand value  $\mu$ ). A PtC score of 0 indicates optimal ordering ( $q = q^*$ ). Referring to Eq. (9), I expect subjects in the RA Algorithm condition to place orders closer to  $q^*$  and exhibit lower PtC bias than those in the control condition.

**H3:** *DMs who observe a risk averse algorithm exhibit lower PtC bias compared to those who do not observe any algorithm in both treatment and post-treatment periods.*

I test the hypotheses in a single-item newsvendor task (study 1). The study was pre-registered on AsPredicted ([https://aspredicted.org/blind.php?x=N9Y\\_BKV](https://aspredicted.org/blind.php?x=N9Y_BKV)).

### 3. Study 1

#### 3.1 Methods

**Subjects.** Subjects were MBA students from a large private university. Study 1 was conducted online, and all participants received a \$5 equivalent Amazon voucher in exchange for participation. The final sample size was 76 (female = 37 %) and the average age in the sample was 30.2 years. In addition to the participation fee, I randomly selected 10% of subjects to win extra Amazon vouchers based on their profits earned in a randomly chosen round of the study. The additional average earnings were \$7 per winning subject.

**Design and Procedure.** Once subjects consented to take part in the study and were briefed of the incentives, they were asked to assume the role of a cart vendor

who bought and sold fish ( $p = \$12, c = \$9$ ). Subjects were informed about the single-item news vendor set-up described earlier in Section 2. The complete set of task instructions used in study 1 is provided in Appendix C2.

Consistent with the setup, the study was comprised of three periods: a baseline period (6 rounds), a treatment period (18 rounds) and a post-treatment period (12 rounds). In every round, subjects observed the demand distribution and decided how much fish to order for that round. The demand distribution for fish was  $\sim U(0 + z_t, 200 + z_t)$ , where  $z_t$  is a random integer between 0 and 40, independently generated for every round (and undisclosed to subjects). I incorporated random parameter  $z_t$  to vary the demand distribution in each round, so that subjects encountered a (slightly) different algorithmic recommendation every time. The demand distribution and actual demand values realized were identical for all subjects in any given round of the study.

At the end of the baseline period, subjects were randomly assigned either to a treatment condition in which they observed recommendations that were generated by a risk-averse algorithm (*RA Algorithm*) or to a control condition, in which they did not encounter any recommendation. Subjects in the treatment condition learned that the risk-averse algorithm tried to strike a balance between maximizing profits and generating more stable returns (i.e., lower variance in profits). I carefully calibrated the level of risk aversion of the algorithm based on previous newsvendor studies (See Appendix C1 for details on algorithm design). Subjects in the treatment condition were free to decide whether or not to follow the algorithm's recommendation. Figure 8 below summarizes the experimental design for study 1.

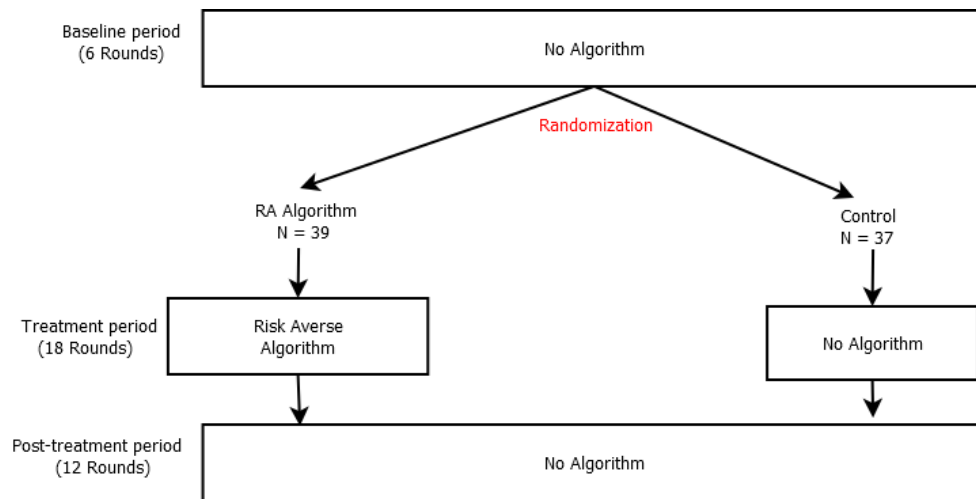


Figure 8. Study Design

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 quantity, the algorithmic recommended order quantity, the actual demand, and whether or not they under- or over-ordered.<sup>17</sup> 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 quantity matched the realized demand value. An example of the feedback screen displayed to subjects is shown in Appendix C3. At the end of the treatment period, I removed the algorithmic aid for all subjects, who then proceeded to complete the final 12 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.

---

<sup>17</sup> The algorithmic recommended order quantity was only displayed to subjects in the RA algorithm condition during the treatment period.

## 3.2 Results

**Ordering behavior.** I operationalized ordering behavior in study 1 as the number of units of fish ordered per round.<sup>18</sup> I analyzed ordering behavior in each of the experimental periods separately. In Figure 9, we observe no discernible differences in ordering behavior between subjects in RA Algorithm and control conditions during the baseline period due to random assignment (also see Table 5, columns 1-2). In Figure 9, we see that subjects in the RA Algorithm condition ordered less than subjects in the control condition during the treatment period (also see Table 5, columns 3-4). For example, subjects in the RA Algorithm condition ordered, on average, 15 fewer units of fish compared to subjects in the control condition during the treatment period. In Figure 9, we still observe differences in ordering behavior among subjects based on their assigned experimental condition during the first half of post-treatment period (Post-I), but not during the second half (Post-II) (also see Table 5, columns 5-8). Additionally, a paired t-test revealed that subjects in the RA Algorithm condition ordered about 15 fewer units of fish in the post-treatment period compared to baseline period ( $t(38) = 3.12, p = 0.003$ ). Thus, study 1 provides support for H1 and H2.

---

<sup>18</sup> We account for variations in ordering behavior resulting from changes in the demand distribution (specifically, the inclusion of the random parameter  $z_t$ ) by incorporating round-fixed effects in all the regression models.

Note: Shaded region signifies  $\pm 1$  standard error




Figure 9. (Mean) Units of Fish Ordered

**Table 5. OLS regression on ordering behavior with round fixed-effects included** (Standard errors are clustered at the subject-level)

		<i>Dependent variable:</i>							
		Units of Fish Ordered							
		Baseline		Treatment		Post-I		Post-II	
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
RA Algorithm		7.407 (7.046)	9.512 (6.113)	-15.588*** (5.107)	-15.459*** (3.936)	-11.226* (6.641)	-12.275** (6.114)	-3.240 (6.965)	-7.283 (6.114)
Age			1.019 (1.289)		0.216 (0.595)		-0.929 (0.840)		-0.547 (0.840)
Gender			11.296* (6.723)		-1.024 (3.743)		-8.991 (6.826)		-10.433 (6.826)
Previous Round Profits			-0.030*** (0.010)		-0.039*** (0.008)		-0.033*** (0.009)		-0.027*** (0.009)
Baseline Order					0.342*** (0.067)		0.205* (0.120)		0.354*** (0.120)
Constant		117.305*** (6.027)	20.702 (40.468)	100.499*** (6.052)	60.251*** (20.523)	99.761*** (6.668)	99.363*** (29.212)	89.808*** (6.161)	86.360*** (29.212)
Observations		456	380	1,368	1,292	456	380	456	380
Adjusted R <sup>2</sup>		0.098	0.111	0.083	0.247	0.014	0.072	0.016	0.090

Note: \*p<0.1; \*\*p<0.05; \*\*\*p<0.01

**Advice Utilization.** During the treatment period, subjects in the RA algorithm condition deviated from algorithmic recommendations by ordering about 24 more units

of fish (on average) than what was recommended. Consistent with previous studies, I computed a weight on advice (WOA) score, which is a measure commonly used to capture the extent of advice utilization (Logg et al., 2019; Soll & Larrick, 2009). I found that subjects in the RA Algorithm condition utilized 74% of the algorithmic recommendation (on average). The WOA score calculation, which involves propensity score matching is detailed in Appendix C4.

**Average Profits.** Subjects in the RA algorithm condition earned higher average profits per round (~\$72) than those in the control condition during the treatment period. However, I did not find any significant difference in average profits earned by subjects between the two conditions during the post-treatment period (See Appendix C4).

**Variability in profits.** I calculated the average standard deviation of profits for each subject across the treatment and post-treatment periods, separately and used this as a measure of variability of profits. From Figure 10 below, we see that subjects in the RA Algorithm condition experienced lower variability in profits compared to those in the control condition during the treatment ( $t(74) = -2.95, p = 0.004$ ) and post-treatment periods ( $t(74) = -1.78, p = 0.079$ ). This finding follows as a consequence of H1 and H2.

Note: Error bars signify  $\pm 1$  standard error




Figure 10. Average Standard Deviation in Profits

**Pull to Center bias.** As described in section 2, a PtC score of 0 indicates optimal ordering ( $q = q^*$ ) and higher values indicate more bias. In Figure 11 below, we observe that the average PtC score decreased significantly for subjects in the RA Algorithm condition from 0.61 in the baseline period to 0.1 ( $p < 0.001$ ) and 0.26 ( $p < 0.001$ ) in the treatment and post-treatment periods, respectively. These results support H3.

Note: Error bars signify  $\pm 1$  standard error

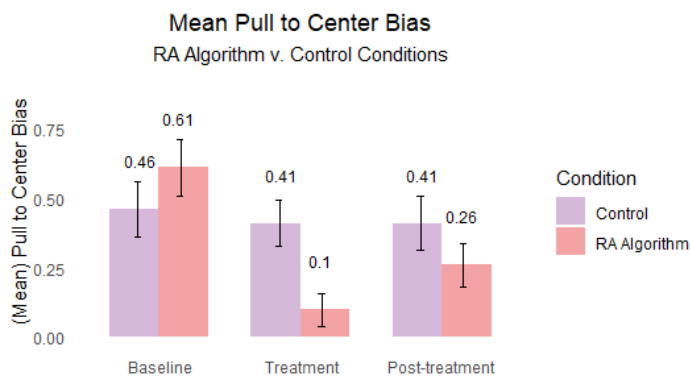


Figure 11. (Mean) Pull to Center Bias

Finally, since there is some debate in the literature regarding the existence of PtC bias as a phenomenon at the individual level (Lau et al., 2014), I conducted an individual, within-subject analysis and confirmed that for a majority of subjects (~70%) in the RA Algorithm condition, the PtC score decreased from baseline to post-treatment period (See Appendix C5). I also checked for other common behavioral biases in the newsvendor setting such as demand chasing and did not find support for it (See Appendix C6).

#### **4. Study 2**

In the previous study, I demonstrated the influence of risk-averse algorithmic anchors in the newsvendor context. The objective of studies 2a and 2b is two-fold. Firstly, in study 2a, I 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. Previous research on algorithm aversion and advice-taking suggests that there could be differences in algorithmic utilization depending on whether the algorithmic assignment is endogenous or exogenous (Dietvorst et al., 2018; Rader et al., 2017). This distinction is also 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 either matching DMs with an algorithm based on a specific risk aversion level or allowing them to choose an algorithm with a particular risk aversion level is common in similar resource allocation contexts (D'Acunto et al., 2019).

Secondly, in study 2b, I aimed to examine whether the *source* of advice (human or algorithmic) influences decision-making in the multi-item newsvendor task. This investigation is important because previous studies on algorithm aversion have

highlighted that people might respond differently to recommendations from humans versus algorithms (Dietvorst et al., 2015). Importantly, examining these factors regarding *autonomy* and *source* helps shed light on the boundary conditions of the observed anchoring effects in newsvendor contexts.

#### **4.1 Study 2a**

The objective of study 2a is to investigate the influence of algorithms in the multi-item newsvendor setting when DMs are free to choose an algorithm with a specific level of risk aversion. This would also help us to understand if DMs have an intrinsic preference for choosing a highly risk-averse algorithmic support. Drawing on previous research on algorithm aversion and advice-taking, I predict that algorithmic utilization would be higher when individuals can select an algorithm (as opposed to one being randomly assigned). Firstly, from the perspective of algorithm aversion, previous research has shown that DMs are more inclined to use algorithms when they are allowed to modify algorithmic recommendations (Dietvorst et al., 2018). I therefore believe that offering individuals the choice to select an algorithm, along with the ability to make judgment-based adjustments, would further enhance the utilization of algorithmic support. Secondly, from the advice-taking literature, we know that individuals place greater confidence in their final estimates when they receive advice that closely aligns with their own initial estimates, as opposed to advice that is randomly chosen (Rader et al., 2017; Yaniv et al., 2009). Therefore, to the extent that individuals select an algorithm based on its proximity to their own risk preference, it is expected that algorithmic utilization will be higher when the choice of algorithm is determined endogenously (by the individual) rather than exogenously (randomly assigned). Study 2a tested this prediction in a pre-registered lab experiment ([https://aspredicted.org/V5K\\_Q2J](https://aspredicted.org/V5K_Q2J)).

## 4.2 Methods

**Subjects.** As in study 1, subjects were Master's degree students. Study 2a was conducted online and subjects earned class credits in return for their participation. In addition to earning extra course credits, I selected 10% of subjects at random and awarded them Amazon vouchers based on profits earned in a randomly chosen round. The average earnings were \$12 per winning subject. The final sample size was 100 (female = 37%) with an average age of 30.5 years.

**Design and Procedure.** Study 2a was designed to closely resemble the multi-item newsvendor study from chapter 2 in all respects, except for one key distinction. At the end of the baseline period, subjects in study 2a were explicitly asked to choose between a highly risk-averse (HRA) algorithm and a risk-neutral (RN) algorithm as their decision support during the treatment period. The purpose of this setup was to compare responses obtained from endogenous choice (participants selecting their preferred algorithm in study2a) versus exogenous allocation (algorithms randomly assigned in chapter 2) by maintaining the overall similarity between the two studies. In fact, these studies were run with two successive cohorts from the same Master's program, and I found no significant differences in demographics or baseline ordering behavior between the two sets of participants. This bolstered the confidence to compare the ordering behavior across both the studies. The methods are reported in detail in Appendix C7.

## 4.3 Results

I find that 35% of subjects chose the HRA algorithm and 65% of subjects chose the RN algorithm and the difference was statistically significant ( $\chi^2(1) = 16.82, p < 0.001$ ).

**Ordering behavior.** I operationalized ordering behavior in study 2a as the number of units of (riskier) product B ordered per round. I analyzed ordering behavior in each of the three experimental periods separately (See Figure 12 below). I find that the ordering behavior of subjects who encounter endogenous risk-neutral (RN) and highly risk-averse (HRA) algorithms is remarkably similar to their respective exogenous counterparts, both in the treatment and post-treatment periods. The regression analysis is reported in Appendix C7.

*Note: Shaded region signifies  $\pm 1$  standard error*

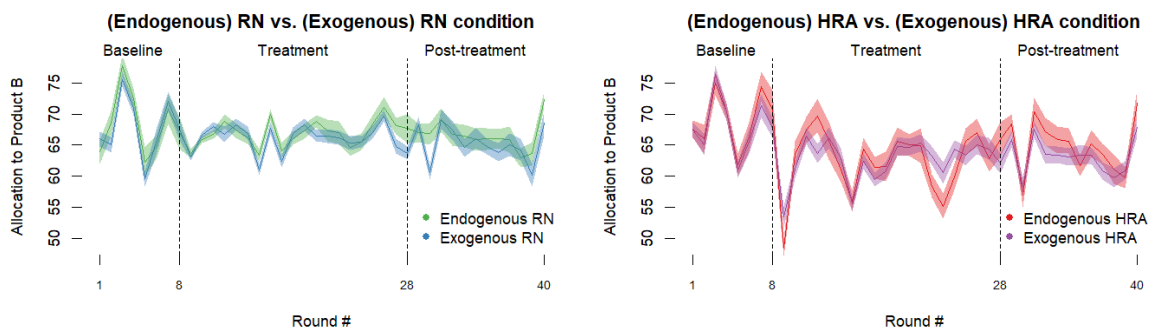


Figure 12. Allocation to Product B (Endogenous v. Exogenous Conditions)

**Advice Utilization.** I used propensity score matching for the weight on advice (WOA) calculation (See Appendix C7). Regardless of whether the algorithm is assigned externally or chosen by the subjects themselves, I find similar levels of advice utilization across both HRA and RN conditions. Additionally, I find that subjects who encounter the RN algorithm utilize the advice about 30% more than those who observe the HRA algorithm ( $\overline{WOA}_{HRA} = 0.52$ ;  $\overline{WOA}_{RN} = 0.83$ ).

**Average Profits and Variability of Profits.** I did not find a significant difference in either average profits or variability of profits during the treatment and post-treatment periods regardless of whether the algorithm (HRA or RN) was externally assigned or chosen by the subjects themselves.

Overall, I find that people utilize algorithmic recommendations to the same extent regardless of whether they have the autonomy to choose the algorithm or not.

#### 4.4 Study 2b

Besides decision autonomy, an important practical consideration for any type of decision support is whether the recommendations come from a human or an algorithm. Thus, the goal of study 2b is to understand whether the source of advice influences the strength and persistence of anchoring effects. Since prior literature has shown that DMs trust algorithms more than humans in tasks that are perceived as more objective, complex and provide opportunities to gain experience, I expect DMs to be more strongly influenced by algorithm (vs. comparable human) recommendations in the multi-item newsvendor setup (Castelo et al., 2019b; Filiz et al., 2021; Mahmud et al., 2022).

Study 2b tested this prediction in a pre-registered lab experiment ([https://aspredicted.org/blind.php?x=V9Y\\_CY4](https://aspredicted.org/blind.php?x=V9Y_CY4))

#### 4.5 Methods

**Subjects.** Subjects were MBA students from a large private university. Study 2b was conducted online and I used a similar incentive scheme employed in previous study 1. The final sample size was 161 (female = 38 %) with an average age of 29.7 years. The average earnings were \$24 per winning subject.

**Design and Procedure.** Study 2b was designed to be similar to the multi-item newsvendor study from chapter 2 in all important respects except for one key distinction.<sup>19</sup> At the end of the baseline period, subjects were randomly assigned to

---

<sup>19</sup> Study 2b had 4 fewer rounds than study 1. Notably, the last 2 rounds in the baseline period and treatment period from study 1 were not present in study 2b. However, we ensured that all 36 rounds in study 2b were identical to 36 rounds in study 1.

one of two treatment conditions in which they observed recommendations that were generated by a highly risk-averse algorithm (*HRA Algorithm*) or a highly risk-averse human participant (*HRA Human*) who had taken part in a similar study previously. The recommendations were identical in terms of average risk aversion level for both the treatment conditions. This design is in line with previous studies on algorithm aversion (Logg et al., 2019) and was pre-registered prior to data collection (For more details regarding the recommendations design, refer to Appendix C8).

#### 4.6 Results.

**Ordering behavior.** I found that subjects in the HRA algorithm condition ordered, on average, about 1.7 fewer units of riskier Product B compared to those in the HRA Human condition during the treatment period (See Figure 13 below). However, there were no significant differences in ordering behavior between the two conditions during the post-treatment period.

*Note: Shaded region signifies  $\pm 1$  standard error*




Figure 13. Allocation to Product B  
(Human v. Algorithm)

**Advice Utilization.** As in study 2a, I used propensity score matching to obtain more accurate initial estimates for the WOA calculation (See Appendix C8). I found

similar levels of advice utilization (WOA  $\approx$  0.5) regardless of the source of advice (human or algorithm).<sup>20</sup>

**Average Profits and Variability of Profits.** I did not find a significant difference in either average profits or variability of profits during the treatment and post-treatment periods irrespective of whether the recommendations came from a human or an algorithm.

Overall, I find that subjects do not exhibit algorithm aversion in the multi-item newsvendor task and utilize HRA human and algorithmic recommendations to a similar extent.

## 5. Discussion

In study 1, I showed that DMs in a single-item newsvendor task ordered less subsequent to observing recommendations from a risk-averse algorithm. Importantly, DMs continued to order less even after the algorithmic recommendation was removed. Thus, in addition to providing DMs risk-neutral algorithmic support (Feng & Gao, 2020; Lee & Siemsen, 2017; Zhang & Siemsen, 2019), feedback and experience (Bolton & Katok, 2008), task training (Bolton et al., 2012) and task decomposition (Lee & Siemsen, 2017b), temporary risk-averse algorithmic support could be a viable complementary strategy to improve newsvendor decisions.

In a follow-up experiment (study 2a), I demonstrated that advice utilization was the same regardless of whether the algorithm was assigned externally or chosen by the subjects themselves in the multi-item newsvendor task. Finally, in study 2b, I

---

<sup>20</sup> The differences in ordering behavior observed between the two treatment conditions (HRA Human and HRA Algorithm) did not directly result in changes in advice utilization. This is because the HRA algorithm and human recommendations were designed to have the same average risk aversion level, but they were not intended to provide identical recommendations in every round.

showed that DMs did not exhibit any algorithm aversion when I compared their responses to highly risk-averse algorithmic recommendations and equivalent human advice. The fact that the main findings remain consistent regardless of critical factors such as *autonomy* and *source of advice* is particularly encouraging for firms contemplating a swift adoption of inventory management software.<sup>21</sup>

Future studies can examine how long the behavioral effects of algorithmic anchors last on order decisions after the recommendation is withdrawn (Kirshner & Moritz, 2022) and additionally test the influence of risk-seeking as well as adaptive algorithms.

I highlight the main results from each of the three chapters in the conclusion section below.

---

<sup>21</sup> Specifically, since algorithmic influence is similar regardless of exogenous or endogenous assignment, organizations could assign an algorithm with a reasonable level of risk-aversion instead of having individual managers choose an algorithm.

## Conclusions

I study individual decision-making in resource allocation tasks when algorithmic support of varying levels of risk is available. Using the literature on anchoring, I predicted that DMs' resource allocation decisions change when they are temporarily exposed to an algorithmic aid. I tested the empirical validity of the claims across a series of experiments in portfolio management and inventory management tasks. In chapter 1, I showed that DMs in a project portfolio management task allocated less resources to the riskier alternative subsequent to observing recommendations from a highly risk-averse algorithm. Importantly, DMs continued to act risk-averse even after the algorithmic recommendation was removed. Interestingly, highly risk-averse algorithmic advice had the strongest effect on changes in resource allocation decisions, despite individuals relying on it the least and hedging away from the advice.

Furthermore, I conceptually replicated the findings from chapter 1 in a more thoughtfully designed inventory management task in chapter 2. This exercise in self-replication helped broaden the applicability of the findings and bolstered the reliability of the original results. Finally, in chapter 3, contrary to expected predictions, I show that critical factors such as decision autonomy and source of advice do not affect the extent the algorithmic utilization in inventory management. Moreover, I document an important use-case of the risk-averse algorithmic support as a debiasing tool in the single-item newsvendor setting.

As a useful reminder, I present the list of main predictions tested in this dissertation below.

1. Highly risk-averse algorithm has the largest impact on resource allocation decisions (Chapters 1 and 2)

2. The impact of an algorithm of any risk-level on resource allocation decisions should depend on whether the algorithm is deliberately chosen by a subject or not (Chapter 3).
3. The impact of a highly-risk averse algorithm on resource allocation decisions should be greater than that of equivalent human advice (Chapter 3).

Without the benefit of hindsight, I would have believed that statements 2 and 3 would have a higher likelihood of being supported than statement 1. However, rather counterintuitively, I found only support for statement 1. Personally, this dissertation stands as a testament to how data can change one's convictions.

The findings have important implications for both decision theory and resource allocation problems. First, unlike recent studies that have demonstrated the persistence of anchoring using “uninformative” anchors (Yoon & Fong, 2019), I show that risk-averse algorithms, which contain informational value, can leave an imprint on the DM's behavior following temporary exposure. This finding is particularly significant given that, in all my studies, DMs are aware of the level of risk aversion of the algorithmic recommendation and are free to decide whether or not to follow to algorithmic recommendation. I believe this is a significant contribution to the decision theoretic work on anchoring.

Second, I believe this work has important domain-specific implication for inventory management. In particular, this dissertation extends the line of research on decision support in newsvendor environments by focusing on temporary risk-averse algorithmic support as a behavioral intervention to inculcate lasting changes in ordering behavior. I believe this marks a notable departure from earlier work that focused primarily on the provision of risk-neutral order recommendations in simple single-item newsvendor tasks (Feng & Gao, 2020; Lee & Siemsen, 2017; Zhang & Siemsen, 2019). In

particular, I show that highly risk-averse algorithms can serve as anchors in cognitively complex inventory problems to sustainably shift order decisions in a risk-averse direction. Such an attitudinal change helps individuals achieve more stable profits over time. The study also provides practical insights for designing algorithmic decision support in inventory management. For instance, the research highlights the critical need for firms to carefully select the level of risk aversion in the algorithm. Lastly, the fact that the main findings remain consistent regardless of critical factors such as *autonomy* and *source of advice* is particularly encouraging for firms contemplating a swift adoption of inventory management software.

The results also have implications for other resource allocation contexts such as portfolio management. For instance, in the domain of finance, subscription-based decision support is becoming increasingly common. My results show that in these contexts, companies could use algorithmic recommendations to nudge employees towards a specific risk-taking behavior. In fact, governments have been trying to increase consumer participation in equity markets (Mehra & Prescott, 1985). In such contexts, temporary algorithmic support could be potentially used to reduce risk aversion and increase investments in equities.

Finally, the experimental design, with three distinct phases (baseline, treatment and post-treatment) can serve as a standard to study the extent of algorithmic support beyond resource allocation problems, such as in sequential search tasks. However, this research does have important limitations that can be addressed in future studies. For instance, the focus was solely on deterministic, risk-averse algorithms, chosen for simplicity and to avoid misunderstandings when explaining risk-seeking algorithms to participants. However, future research can examine the role of adaptive algorithms (whose risk attitude can change based on the DM's choices) and risk-seeking

algorithms. Additionally, due to practical constraints with a student population and the need to complete the experiment within a reasonable timeframe, I limited the post-treatment phase to a maximum of twelve rounds. Future work could investigate the persistence of algorithmic influence over longer periods. Lastly, I confined the study to scenarios involving gains to prevent the confounding effects of risk and loss aversion. This restriction could be overcome in future research by carefully introducing mixed gambles that include both gains and losses.

Overall, this dissertation offers insights into the important role algorithmic recommendations can play in influencing decision-making under risk. This research demonstrates that in resource allocation tasks that involve critical trade-offs between risk and return, algorithms can serve as a useful tool for purposefully inducing lasting changes in human behavior.

## Conclusiones

Estudiamos la toma de decisiones individuales en tareas de asignación de recursos cuando se dispone de apoyo algorítmico con diferentes niveles de riesgo. Utilizando la literatura sobre anclaje, predijimos que las decisiones de asignación de recursos de los DM (tomadores de decisiones) cambian cuando están expuestos temporalmente a una ayuda algorítmica. Probamos la validez empírica de nuestras afirmaciones a través de una serie de experimentos en tareas de gestión de carteras y de inventarios. En el capítulo 1, mostramos que los DM en una tarea de gestión de carteras de proyectos asignaron menos recursos a la alternativa más arriesgada después de observar recomendaciones de un algoritmo altamente reacio al riesgo. Importante, los DM continuaron actuando de manera reacia al riesgo incluso después de que se retirara la recomendación algorítmica. Curiosamente, el consejo algorítmico altamente reacio al riesgo tuvo el efecto más fuerte en los cambios en las decisiones de asignación de recursos, a pesar de que los individuos dependían menos de él y se alejaban del consejo.

Además, replicamos conceptualmente los hallazgos del capítulo 1 en una tarea de inventario multi-elemento más cuidadosamente diseñada en el capítulo 2. Finalmente, en el capítulo 3, contrariamente a las predicciones esperadas, mostramos que factores críticos como la autonomía de decisión y la fuente de consejo no afectan el grado de utilización algorítmica en la gestión de inventarios. Además, documentamos un caso de uso importante del apoyo algorítmico reacio al riesgo como una herramienta para eliminar sesgos en el contexto de un nuevo proveedor de un solo artículo. Estos hallazgos tienen implicaciones importantes tanto para la teoría de la decisión como para las operaciones conductuales. Como un recordatorio útil,

presento a continuación la lista de las principales predicciones probadas en esta disertación.

1. El algoritmo con mayor aversión al riesgo tiene el mayor impacto en las decisiones de asignación de recursos (Capítulos 1 y 2).

2. El impacto de un algoritmo de cualquier nivel de riesgo en las decisiones de asignación de recursos debería depender de si el algoritmo es elegido deliberadamente por un sujeto o no (Capítulo 3).

3. El impacto de un algoritmo con alta aversión al riesgo en las decisiones de asignación de recursos debería ser mayor que el de un consejo humano equivalente (Capítulo 3).

Sin el beneficio de la retrospectiva, habría creído que las declaraciones 2 y 3 tendrían una mayor probabilidad de ser respaldadas que la declaración 1. Sin embargo, de manera bastante contraintuitiva, solo encontré apoyo para la declaración 1. Personalmente, esta disertación se erige como un testimonio de cómo los datos pueden cambiar las convicciones de uno.

Nuestros hallazgos tienen implicaciones importantes tanto para la teoría de la decisión como para los problemas de asignación de recursos. Primero, a diferencia de estudios recientes que han demostrado la persistencia del anclaje utilizando anclas "no informativas" (Yoon & Fong, 2019), mostramos que los algoritmos reacios al riesgo, que contienen valor informativo, pueden dejar una huella en el comportamiento del DM después de una exposición temporal. Este hallazgo es particularmente significativo dado que, en todos nuestros estudios, los DM son conscientes del nivel de aversión al riesgo de la recomendación algorítmica y son libres de decidir si seguir o no la recomendación algorítmica. Creemos que esto es una contribución significativa al trabajo teórico de decisión sobre el anclaje.

Segundo, creemos que nuestro trabajo tiene importantes implicaciones específicas del dominio para la gestión de inventarios. En particular, esta disertación extiende la línea de trabajo sobre apoyo en la toma de decisiones en entornos de nuevos proveedores al enfocarse en el apoyo algorítmico temporal reactivo al riesgo como una intervención conductual para inculcar cambios duraderos en el comportamiento de pedidos. Creemos que esto marca una notable salida del trabajo anterior que se centró principalmente en la provisión de recomendaciones de pedido neutrales al riesgo en tareas simples de nuevos proveedores de un solo artículo (Feng & Gao, 2020; Lee & Siemsen, 2017; Zhang & Siemsen, 2019). En particular, demostramos que los algoritmos altamente aversos al riesgo pueden servir como anclas en problemas de inventario cognitivamente complejos para cambiar de manera sostenible las decisiones de pedido en una dirección aversa al riesgo. Tal cambio actitudinal ayuda a las personas a lograr beneficios más estables a lo largo del tiempo. Nuestro estudio también proporciona información práctica para diseñar apoyo a la decisión algorítmica en la gestión de inventarios. Por ejemplo, nuestra investigación destaca la necesidad crítica de que las empresas seleccionen cuidadosamente el nivel de aversión al riesgo en el algoritmo. Por último, el hecho de que nuestros principales hallazgos permanezcan consistentes independientemente de factores críticos como la autonomía y la fuente de consejo es particularmente alentador para las empresas que contemplan una adopción rápida del software de gestión de inventarios.

Nuestros resultados también tienen implicaciones para otros contextos de asignación de recursos, como la gestión de carteras. Por ejemplo, en el dominio de las finanzas, el apoyo a la decisión basado en suscripciones se está volviendo cada vez más común. Nuestros resultados muestran que, en estos contextos, las empresas

podrían usar recomendaciones algorítmicas para impulsar un comportamiento específico de toma de riesgos. De hecho, los gobiernos han intentado aumentar la participación de los consumidores en los mercados de acciones (Mehra & Prescott, 1985). En tales contextos, el apoyo algorítmico temporal podría utilizarse potencialmente para reducir la aversión al riesgo y aumentar las inversiones en acciones.

Finalmente, nuestro diseño experimental, con tres fases distintas (línea base, tratamiento y post-tratamiento), puede servir como un estándar para estudiar el alcance del apoyo algorítmico más allá de los problemas de asignación de recursos, como en tareas de búsqueda secuencial. Sin embargo, esta investigación tiene limitaciones importantes que pueden ser abordadas en estudios futuros. Por ejemplo, nuestro enfoque se centró únicamente en algoritmos deterministas y aversos al riesgo, elegidos por simplicidad y para evitar malentendidos al explicar algoritmos que buscan el riesgo a los participantes. Sin embargo, investigaciones futuras pueden examinar el papel de algoritmos adaptativos (cuya actitud hacia el riesgo puede cambiar basada en las elecciones del DM) y algoritmos que buscan el riesgo. Además, debido a limitaciones prácticas con una población estudiantil y la necesidad de completar el experimento en un plazo razonable, limitamos la fase de post-tratamiento a un máximo de doce rondas. Trabajos futuros podrían investigar la persistencia de la influencia algorítmica durante períodos más largos. Por último, confinamos nuestro estudio a escenarios que involucran ganancias para prevenir los efectos confusos de la aversión al riesgo y la pérdida. Esta restricción podría superarse en investigaciones futuras introduciendo cuidadosamente apuestas mixtas que incluyan tanto ganancias como pérdidas.

En general, esta disertación ofrece perspectivas sobre el importante papel que las recomendaciones algorítmicas pueden desempeñar en la influencia de la toma de decisiones bajo riesgo. Nuestra investigación demuestra que en tareas de asignación de recursos que involucran compromisos críticos entre riesgo y retorno, los algoritmos pueden servir como una herramienta útil para inducir de manera intencional cambios duraderos en el comportamiento humano.

## Appendix A

- A1. Proof
- A2. Task instructions
- A3. Sample stimuli and participant feedback
- A4. Algorithm Design
- A5. Risk taking behavior
- A6. Deviation from algorithmic recommendation
- A7. Earnings

### Appendix A1. Proof

Consider a DM facing the task of investing her endowment  $W$  in two prospects: a more-risky (MR) and a less-risky (LR) alternative. I assume that MR alternative has a higher mean ( $\mu_1$ ) and a higher coefficient of variation ( $\sigma_1/\mu_1$ ) than LR with mean  $\mu_2$  and coefficient of variation ( $\sigma_2/\mu_2$ ). The returns are assumed to be correlated with a value  $\rho$ . Moreover, consider that the DM has an exponential utility function  $U(W) = -e^{-\lambda W}$ , where  $\lambda > 0$  measures the degree of risk aversion. The utility function is increasing and concave since  $U'(W) = \lambda e^{-\lambda W} > 0$  and  $U''(W) = -\lambda^2 e^{-\lambda W} < 0$ . The Pratt-Arrow measure of absolute risk aversion is constant, i.e.,  $-\frac{U''(W)}{U'(W)} = \lambda$ . Moreover, assume that endowment  $W$  is distributed normally with mean  $\mu$  and standard deviation  $\sigma$  such that the density of  $W$  is given by  $f(W) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(W-\mu)^2/2\sigma^2}$ . Therefore, the expected utility  $EU(W)$  can be described as

$$EU(W) = -\frac{1}{\sigma\sqrt{2\pi}} \int e^{-\frac{(W-\mu)^2}{2\sigma^2}} (-e^{-\lambda W}) dW, \quad (11)$$

which gives  $EU(W) = -e^{-\lambda(\mu - \frac{\lambda\sigma^2}{2})}$ .

Hence, the objective of the DM is to maximize  $\mu - \frac{\lambda\sigma^2}{2}$ , where  $\mu$  denotes the mean return of the portfolio and  $\sigma$  is the standard deviation of the portfolio.<sup>22</sup>

---

<sup>22</sup> Similar result is also obtained with a quadratic utility function (see Sargent, 1981 for detailed discussion).

Let  $\alpha$  and  $(1 - \alpha)$  be the fraction of the endowment allocated to alternatives MR and LR, respectively. The objective function  $F$  is then given by

$$F = \max_{\alpha} (\alpha\mu_1 + (1 - \alpha)\mu_2 - \frac{\lambda}{2} (\alpha^2 \sigma_1^2 + (1 - \alpha)^2 \sigma_2^2 + 2\alpha(1 - \alpha)\rho\sigma_1\sigma_2)). \quad (12)$$

**Lemma 1.** *The optimal allocation  $\alpha^*$  to MR alternative decreases with increasing risk aversion  $\lambda$ .*

**Proof.**

$$f = \max_{\alpha} (\alpha\mu_1 + (1 - \alpha)\mu_2 - \frac{\lambda}{2} (\alpha^2 \sigma_1^2 + (1 - \alpha)^2 \sigma_2^2 + 2\alpha(1 - \alpha)\rho\sigma_1\sigma_2)). \quad (13)$$

Taking the first derivative of  $f$  I get

$$\frac{\partial f}{\partial \alpha} = \mu_1 - \mu_2 - \frac{\lambda}{2} [2\alpha\sigma_1^2 - 2(1 - \alpha)\sigma_2^2 + 2(1 - \alpha)\sigma_2^2 + 2(1 - \alpha)\rho\sigma_1\sigma_2 - 2\alpha\rho\sigma_1\sigma_2]. \quad (14)$$

The second derivative of  $f$  is given by

$$\frac{\partial^2 f}{\partial \alpha^2} = -\frac{\lambda}{2} [2\sigma_1^2 + 2\sigma_2^2 - 4\rho\sigma_1\sigma_2] < 0. \quad (15)$$

The second derivative is negative as  $\sigma_1 > \sigma_2$ . Equating the first derivative to zero, I get

$$\alpha^* = \frac{\frac{2}{\lambda} (\mu_1 - \mu_2) + 2\sigma_2^2 - 2\rho\sigma_1\sigma_2}{2\sigma_1^2 + 2\sigma_2^2 - 4\rho\sigma_1\sigma_2}. \quad (16)$$

In the above expression I can observe that  $\frac{\partial \alpha^*}{\partial \lambda} < 0$  as  $\mu_1 - \mu_2 > 0$ . Therefore, the optimal allocation  $\alpha^*$  to MR alternative decreases with increasing  $\lambda$ .

## Appendix A2. Task instructions

I report below the screenshots of all task instructions from the study as they would appear on a participant's computer screen until round 1. I additionally report the screenshots when assigning participants to different treatment conditions (HRA, SRA and RN) as well as for the first round in the treatment period (round 9). Screenshots of the remaining parts of the study are available upon request.

### **General Structure**

This study involves a resource allocation game that consists of **40 rounds**. In each round, as a Manager you will be required to allocate your company's resources between two projects.

**The goal of the game is to maximize the overall return.**

The expected time to complete the study is about 30-45 minutes.

Please complete the entire study in one sitting.



### Task Overview

You will make a total of 40 resource allocation decisions (or 40 rounds) during the course of the study.

The 40 rounds in the study are divided into 5 sessions of 8 rounds each. At the beginning of each session, you will be endowed with a budget of \$10. In the subsequent rounds of that session, your endowment will change depending on your allocation decision and the performance of the two projects during each round.

In each round, you will be asked to choose the percentage of your current endowment that you would like to allocate between 2 options: a more-risky project (**MR**) and a less-risky project (**LR**).

The below table shows the past mean returns and standard deviation of each project. Feel free to note it down somewhere.

	Mean Return	Standard Deviation
More-Risky (MR) project	9.7 %	4 %
Less-Risky (LR) project	7.9 %	2.5 %

**Note:** MR project is riskier than LR project (since it has a higher standard deviation).

MR project also has a higher mean return than LR project.

At the end of the each round, you will see the performance of your portfolio and the accumulated wealth.



### Incentives

All participants will earn class extra credits.

In addition, you have the chance to win an Amazon voucher. We will randomly select 10 participants to win the voucher. We will then randomly choose one of the five sessions and award those 10 lucky participants the value of their total wealth accumulated in that session as an Amazon voucher. A participant can earn up to a maximum of \$25 (approx.)

Therefore, please make your decisions carefully in every round.



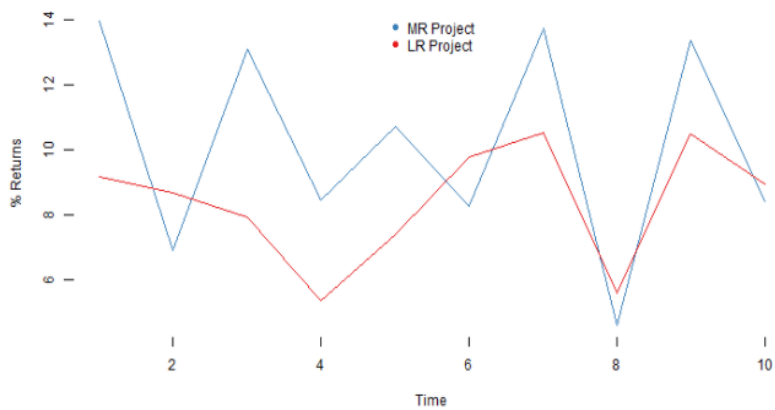
You are now ready to start Session 1 with an endowment of \$10.

Let's begin...



### Round 1 of 40

The graph below shows the return of the projects for the last 10 periods.



Please allocate your current wealth to the two projects based on your forecast for the next period.

For example, if you want to invest 25% of your wealth in one option, please write 25.

Total must sum to 100.

More-Risky (MR) Project Allocation (%)

Less-Risky (LR) Project Allocation (%)

Total

## HRA condition

Now, in each round there will be a **decision support model** available, which recommends an optimal allocation of your resources. The model is based on a **complex, mathematical algorithm**. Importantly, the algorithm is **highly risk-averse** in optimizing the portfolio.

**Degree of risk aversion** typically ranges on a scale from 0-5, with higher values indicating **more conservative** (risk-averse) investment. The algorithm is **highly risk-averse** and uses a **degree of risk-aversion of value 5**.

**Note that whether you follow the algorithmic suggestion or not is entirely up to you.**




The **algorithm** assigned to you is

- Highly Risk-Averse
- Slightly Risk-Averse
- Not at all Risk-Averse



Round 9 of 40

The graph below shows the the return of the projects for the last 10 periods.



The algorithm predicts the following returns for the next period.

More-Risky Project Forecast	Less-Risky Project Forecast
8.36 %	6.28 %

Based on the forecast, the algorithm suggests the following wealth allocation, which is **highly risk-averse**.

More-Risky Project Allocation (%)	Less-Risky Project Allocation (%)
31	69

Please allocate your current wealth to the two projects based on your forecast for the next period.

For example, if you want to invest 25% of your wealth in one option, please write 25. Total must sum to 100.

More-Risky (MR) Project Allocation (%)	<input type="text" value="0"/>
Less-Risky (LR) Project Allocation (%)	<input type="text" value="0"/>
Total	<input type="text" value="0"/>

### SRA condition

Now, in each round there will be a **decision support model** available, which recommends an optimal allocation of your resources. The model is based on a **complex, mathematical algorithm**. Importantly, the algorithm is **slightly risk-averse** in optimizing the portfolio.

**Degree of risk aversion** typically ranges on a scale from 0-5, with higher values indicating **more conservative** (risk-averse) investment. The algorithm is **slightly risk-averse** and uses a **degree of risk-aversion of value 0.5**.

**Note that whether you follow the algorithmic suggestion or not is entirely up to you.**




The **algorithm** assigned to you is

- Highly Risk-Averse
- Slightly Risk-Averse
- Not at all Risk-Averse



**Round 9 of 40**

The graph below shows the the return of the projects for the last 10 periods.



The **algorithm** predicts the following returns for the next period.

More-Risky Project Forecast	Less-Risky Project Forecast
8.36 %	6.28 %

Based on the forecast, the **algorithm** suggests the following wealth allocation, which is **slightly risk-averse**.

More-Risky Project Allocation (%)	Less-Risky Project Allocation (%)
46	54

Please allocate your current wealth to the two projects based on your forecast for the next period.

For example, if you want to invest 25% of your wealth in one option, please write 25.

Total must sum to 100.

More-Risky (MR) Project Allocation (%)	<input type="text" value="0"/>
Less-Risky (LR) Project Allocation (%)	<input type="text" value="0"/>
<b>Total</b>	<input type="text" value="0"/>

## RN condition

Now, in each round there will be a **decision support model** available, which recommends an optimal allocation of your resources. The model is based on a **complex, mathematical algorithm**. Importantly, the algorithm is **risk-neutral** in optimizing the portfolio.

**Risk-neutral** algorithm suggests an allocation which maximises returns **without** any consideration for risk.

**Note that whether you follow the algorithmic suggestion or not is entirely up to you.**



The **algorithm** assigned to you is

- Highly Risk-Averse
- Slightly Risk-Averse
- Risk-Neutral



Round 9 of 40

The graph below shows the the return of the projects for the last 10 periods.

Time	MR Project (%)	LR Project (%)
1	8.5	9.0
2	7.0	8.0
3	5.5	7.5
4	7.5	8.0
5	12.0	7.8
6	10.0	7.5
7	7.5	7.8
8	12.0	7.8
9	7.0	7.5
10	8.0	9.5

The algorithm predicts the following returns for the next period.

More-Risky Project Forecast	Less-Risky Project Forecast
8.36 %	6.28 %

Based on the forecast, the algorithm suggests the following wealth allocation, which is risk-neutral.

More-Risky Project Allocation (%)	Less-Risky Project Allocation (%)
100	0

Please allocate your current wealth to the two projects based on your forecast for the next period.  
 For example, if you want to invest 25% of your wealth in one option, please write 25.  
 Total must sum to 100.

More-Risky (MR) Project Allocation (%)

Less-Risky (LR) Project Allocation (%)

Total

### Appendix A3. Sample stimuli and participant feedback

The performance of each project during the past 10 rounds is shown in the form of a graph in every round (e.g., see Figure 14). Once a participant makes a resource allocation decision, a feedback screen is displayed as shown in Figure 15.

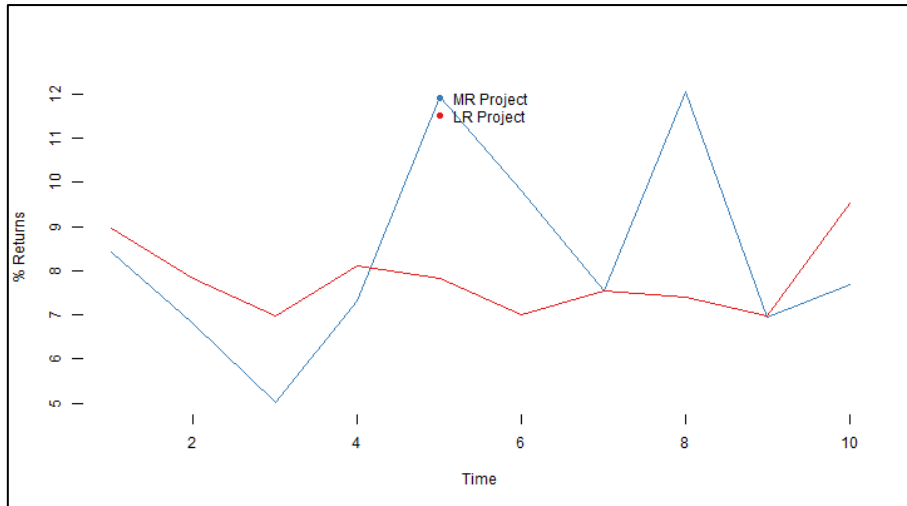


Figure 14. Past projects' returns shown to participants in the study

Return of the algorithmic suggestion	7.05 %
Highest possible return you could have earned	7.44 %
Your return	6.69 %
Your current wealth (in dollars)	10.67

Figure 15. Feedback screen displayed to a hypothetical participant in the study

#### Appendix A4. Algorithm Design

I outline below the process for generating the algorithmic forecast of the projects' returns and subsequently the algorithmic suggestion of optimal resource allocation. The algorithmic forecast generated is identical for participants in all treatment conditions (HRA, SRA and RN), whereas the algorithmic suggestion varies based on the assigned treatment condition.

**Algorithmic forecast.** I developed two separate linear regression models to forecast MR and LR project returns in every round. Projects returns were assumed to correlate with an underlying set of independent variables  $X_1$ ,  $X_2$  and  $X_3$  and I used the Cholesky-Decomposition (CD) method to generate correlated random variables (Lurie & Goldberg, 1998). CD in essence works as follows: The inputs needed to perform CD are the initial set of independent random variables (say, X and Y) and a desired correlation value "r". The output of CD are two transformed variables (say, New X and

New Y) which will be correlated with the specified correlation value “r”. The process of CD can be extended to generate multiple correlated random variables. In the study, I used CD to generate MR and LR project returns, which were assumed to be correlated with the underlying set of independent variables  $X_1$ ,  $X_2$  and  $X_3$ . I confirmed that both OLS models exhibited a comparable level of accuracy<sup>23</sup>.

**Algorithmic suggestions.** To generate algorithmic suggestions associated with different risk preferences, I relied on the mean-variance framework. The objective function for a decision-maker (DM) takes the following functional form:

$$E(r) - \lambda * Var(r) , \quad (17)$$

where  $E(r)$  denotes the expected returns (i.e., algorithmic forecast),  $Var(r)$  refers to the variance in returns and  $\lambda$  represents a parameter describing the DM’s degree of risk aversion. Higher values of  $\lambda$  indicate more risk version. The value of  $\lambda$  is 0 for a risk neutral DM. I chose  $\lambda$  values of 0.5 and 4 to model the SRA and HRA algorithms respectively<sup>24</sup>. For each treatment condition, I ran a Monte Carlo simulation to determine the optimal wealth allocation between MR and LR projects that maximized the objective function in Eq. (17) in every round.

## Appendix A5. Resource allocation decisions

It is possible that decisions in the post-treatment period is affected by the earnings (i.e., returns or profits) gained during treatment period. For instance, (Thaler & Johnson, 1990) report changes in risk taking behavior after people experience gains. To account for this explanation, I control for the average earnings gained by participants during the treatment period and the regression results are reported below in Table 6. I again observe that participants in HRA condition are significantly more

---

<sup>23</sup> A paired t-test revealed no significant differences in the mean absolute error in predictions between the two regression models during the treatment period ( $t(23) = -1.67$ ,  $p = 0.11$ ). Mean absolute error in predictions were calculated as the absolute value of the difference between actual (MR or LR) and forecasted (MR or LR) project returns.

<sup>24</sup> For the chosen values of  $\lambda$ , we confirmed that the HRA algorithmic suggestion was significantly more risk-averse than the SRA algorithmic suggestion. A paired t-test revealed significant mean differences in allocation to the LR project between SRA and HRA algorithms during the treatment period ( $t(23) = -8.18$ ,  $p < 0.001$ ).

risk-averse compared to those in the control condition during the post-treatment period.

**Table 6. OLS regression on risk taking behavior with round-fixed effects included.**

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

	<i>Dependent variable:</i>	
	Endowment Allocated to MR Project Post-Treatment	
	(1)	(2)
Highly Risk-averse	-8.787** (3.617)	-6.656** (3.173)
Risk Neutral	-2.277 (3.367)	-5.885* (3.288)
Slightly Risk-averse	-4.456 (3.237)	-8.258*** (3.096)
Age		-0.310 (0.330)
Gender		-0.961 (2.183)
Average Treatment Returns		23.196*** (5.000)
Previous Round Returns		3.465* (2.041)
Baseline Allocation		0.442*** (0.091)
Constant	75.316*** (2.680)	-188.405*** (44.350)
Observations	1,576	1,379
Adjusted R <sup>2</sup>	0.030	0.129

*Note:* \*p<0.1; \*\*p<0.05; \*\*\*p<0.01

## Appendix A6. Deviation from algorithmic recommendations

I report regression results for the deviation from the algorithmic recommendation as the dependent variable in Table 7. As a robustness check, I also ran regression analyses using the normalized deviation from algorithmic recommendation and the results are qualitatively similar (see Table 8). The absolute and normalized deviations were computed as follows:

$$\text{Deviation (absolute)} = \text{Subject's Allocation} - \text{Algorithmic Recommendation} \quad (18)$$

$$\text{Deviation (normalized)} = \frac{\text{Subject's Allocation} - \text{Algorithmic Recommendation}}{\text{Algorithmic Recommendation}} \quad (19)$$

**Table 7. OLS regression on deviation from algorithm with round-fixed effects included.** Standard errors are clustered at the subject-level and reported in parentheses.

	<i>Dependent variable:</i>	
	(Absolute) Deviation from algorithm Treatment	
	(1)	(2)
Risk Neutral	-25.111*** (2.460)	
Highly Risk-averse	22.365*** (2.394)	49.261*** (2.293)
Slightly Risk-averse	8.247*** (2.329)	33.181*** (2.175)
Age		-0.123 (0.248)
Gender		1.196 (1.840)
Trust in Algorithm		0.881 (0.649)
Previous Round Returns		3.750*** (0.850)
Baseline Allocation		0.260*** (0.058)
Constant		-73.070*** (10.258)
Observations	3,504	3,358
Adjusted R <sup>2</sup>	0.350	0.371

Note: \*p<0.1; \*\*p<0.05; \*\*\*p<0.01

**Table 8. OLS regression on deviation from algorithm with round-fixed effects included.** Standard errors are clustered at the subject-level and reported in parentheses.

	<i>Dependent variable:</i>	
	(Normalized) Deviation from algorithm Treatment	
	(1)	(2)
Risk Neutral	-25.665*** (3.057)	
Highly Risk-averse	27.043*** (3.548)	54.590*** (3.014)
Slightly Risk-averse	17.625*** (3.913)	42.668*** (3.253)
Age		-0.086 (0.330)
Gender		2.281 (2.650)
Trust in Algorithm		0.715 (0.879)
Previous Round Returns		5.331*** (1.391)
Baseline Allocation		0.464*** (0.090)
Constant		-92.187*** (14.993)
Observations	3,504	3,358
Adjusted R <sup>2</sup>	0.244	0.264

*Note:* \*p<0.1; \*\*p<0.05; \*\*\*p<0.01  
The intercept has been forced to 0 in column 1

## Appendix A7. Earnings

Table 9 below shows subjects' mean returns earned (per round) with standard deviation in parentheses in each of the three experimental periods.

**Table 9. Returns earned in the study**

	<b>Baseline</b>	<b>Treatment</b>	<b>Post-treatment</b>
<b>Control</b>	8.17 (0.24)	8.86 (0.24)	8.20 (0.13)
<b>HRA</b>	8.16 (0.23)	8.83 (0.23)	8.14 (0.14)
<b>SRA</b>	8.21 (0.26)	9.04 (0.25)	8.17 (0.16)
<b>RN</b>	8.17 (0.20)	9.04 (0.22)	8.17 (0.14)

## Appendix B

- B1. Algorithm design
- B2. Task instructions
- B3. Participant feedback
- B4. Deviation and average profits
- B5. Alternate explanations

### Appendix B1. Algorithm design

I design the algorithm(s) such that it tries to maximize an objective function  $F$ , which can be expressed as follows:

$$F = E(\pi_t) - \lambda * V(\pi_t) \quad (20)$$

In Eq. (20),  $E(\pi_t)$  is the expected profits,  $V(\pi_t)$  represents the variance in profits and  $\lambda$  captures the degree of risk aversion of the algorithm. The intuition behind Eq. (20) is that an increase in risk aversion is modeled by an increasing distaste for payoff variance. Using simulation, for the study parameters I used, I show that *higher values of risk aversion  $\lambda$  lead to less resource allocation to riskier product B*. This result is shown in Figure 16.




Figure 16. Allocation to Product B v. Risk Aversion ( $\lambda$ )

*Note: The simulation result is shown for a particular decision round with a specific demand distribution ( $D_{At} = U(15,50), D_{Bt} = U(0,125)$ ). However, I confirmed that the allocation to product B decreases with increasing risk aversion  $\lambda$  for all demand distributions I used in the study.*

## Appendix B2. Task instructions

I report below the screenshots of all task instructions from the study as they would appear on a participant's computer screen until round 1. I additionally report the screenshots when assigning participants to different treatment conditions (HRA, SRA and RN) as well as for the first round in the treatment period. Screenshots of the remaining parts of the study are available upon request.

### General Structure

This study involves an inventory ordering game that consists of 40 rounds.

The expected time to complete the study is about 30-45 minutes.

Please complete the entire study in one sitting.



### Task Overview

You are a **Manager** of a retail store and your store makes profits by selling two kinds of **perishable, non-substitutable products - Product A and Product B.**



You will be making inventory decisions over 40 rounds.  
The 40 rounds in the study are further divided into 3 sessions.

In each round, **you are solely responsible for deciding how many quantities of each product type you would like to order from a supplier.**

The store operations process is described in the next screen.



## Store Operations



**Please read the operations carefully.**

At the beginning of each round, you will decide to stock some quantities of Product A and Product B that you would like to sell to customers in the same round.

The unit prices are shown in the Table below. They remain fixed for all the rounds and are identical for both Products A and B.

Product Type	Unit Selling Price (p)	Unit Cost (c)	Unit Profit (p - c)
A	\$12	\$3	\$9
B	\$12	\$3	\$9



## Under- and Over-Ordering



**Over-ordering:** In every round, if you end up **ordering more units than the actual customer demand** in the round, the excess units will perish and would have be discarded - i.e. you would make a loss of \$3 for every excess unit for either product A or B.

**Under-ordering:** However, if you end up **ordering less than the actual customer demand** in the round, you will miss out on potential sales opportunity - i.e. you would lose out on making a profit of \$9 for every unit of excess demand for either product A or B.

**You make maximum profits when your order quantities for Products A and B match the customer demands for these products.**

An example in a round where Product A is under-ordered and Product B is over-ordered is shown in the table below:

( $p = \$12$ ,  $c = \$3$ ).

Product Type	Your Order Quantity	Actual Demand		Your Profits	Max. Possible Profits
A	20	< 30	⇒ Under-ordering	\$180 (20*9)	\$ 270 (30*9)
B	80	> 50	⇒ Over-ordering	\$ 360 (50*9 - 30 *3)	\$ 450 (50*9)
<b>Total Profits (A + B)</b>				<b>\$540</b>	<b>\$720</b>



### Customer Demand

The demand for Products A and B is uncertain and varies in every round of the study. Importantly, **the actual demand for both Products A and B will only be known at the end of the round** (after the order quantities are decided).

However, you will be given information on the demand distribution for both products A and B at the beginning of every round:

Product Type	Range of Possible Demand	Average Demand	Standard Deviation of Demand	Riskiness
Product A	15 - 50	32.5	10.1	Low
Product B	0 - 125	62.5	36.1	High

**The following set of instructions is applicable for all the 40 rounds. Please read them extremely carefully before proceeding with the study.**

1. The actual demand values will be randomly drawn from the range of possible demand values described given to you in a particular round. **Every integer value of demand in the ranges described is equally likely.** For example, the actual demand for Product A can be any integer value in the range of 15 to 50 with equal probability. Similarly, the actual demand for Product B can be any integer value in the range of 0 to 125 with equal probability.

2. Product B has a **higher average demand** than Product A; Product B is also **riskier** (has a higher standard deviation of demand) than Product A.

3. In any round, the actual demand for Product A **does not affect** the demand for Product B (and vice-versa). The actual demand for both Products A and B in any round **is unaffected** by their demand in other rounds.

4. In every round, you will have to order a total quantity of **100 units** (of Products A and B combined). This restriction is in place because of the limited shelf space in your store.

5. At the end of each round, you will see your order quantities, the actual demand values, store profits, maximum possible profits and whether or not you under- or over-ordered.



### **Incentives**

All participants will earn class extra credits.

In addition, you have the chance to win an Amazon voucher. We will randomly select 10 participants to win the voucher. We will then randomly choose 1 of the 40 rounds and award those 10 lucky participants an Amazon voucher whose worth is proportional to the profits generated in that round. You can earn a maximum of €27.

**Therefore, please make your decisions carefully in every round!**



**Let's begin Session 1 ...**



### Round 1

The table below shows the demand distribution for Products A and B in this round.

Product Type	Range of Possible Demand Values
Product A	15 - 50
Product B	0 - 125

Please enter the quantities of Products A and B that you would like to order in this round.  
Total must sum to 100.

Product A	<input type="text" value="0"/>
Product B	<input type="text" value="0"/>
Total	<input type="text" value="0"/>



## HRA condition

Now, in each round there will be a **decision support model** available, which recommends optimal order quantities of Products A and B. The model is based on a **complex, mathematical algorithm**. Importantly, the algorithm is **highly risk-averse** in optimizing the order mix. By that, we mean the algorithm has a general tendency to prefer allocating **more resources** to the **less risky product A**.

**Degree of risk aversion** typically ranges on a scale from 1- 4, with higher values indicating **more conservative** (risk-averse) decisions. The algorithm is **highly risk-averse** and uses a **degree of risk-aversion of value 4**.

**Note that whether you follow the algorithmic suggestion or not is entirely up to you.**



The **algorithm** assigned to you is

- Highly Risk-Averse
- Slightly Risk-Averse
- Not at all Risk-Averse



### Round 9

The table below shows the demand distribution for Products A and B in this round.

Product Type	Range of Possible Demand Values
Product A	23 - 58
Product B	8 - 133

The **algorithm** suggests the following order quantities for Products A and B, which is **highly risk-averse**.

Product A	Product B
58	42

Please enter the quantities of Products A and B that you would like to order in this round.  
Total must sum to 100.

Product A	<input type="text" value="0"/>
Product B	<input type="text" value="0"/>
Total	<input type="text" value="0"/>



## SRA condition

Now, in each round there will be a **decision support model** available, which recommends optimal order quantities of Products A and B. The model is based on a **complex, mathematical algorithm**. Importantly, the algorithm is **slightly risk-averse** in optimizing the order mix. By that, we mean the algorithm has a general tendency to prefer allocating **slightly more resources** to the **less risky product A**.

**Degree of risk aversion** typically ranges on a scale from 1- 4, with higher values indicating **more conservative** (risk-averse) decisions. The algorithm is **slightly risk-averse** and uses a **degree of risk-aversion of value 1**.

**Note that whether you follow the algorithmic suggestion or not is entirely up to you.**



The **algorithm** assigned to you is

- Highly Risk-Averse
- Slightly Risk-Averse
- Not at all Risk-Averse



### Round 9

The table below shows the demand distribution for Products A and B in this round.

Product Type	Range of Possible Demand Values
Product A	23 - 58
Product B	8 - 133

The **algorithm** suggests the following order quantities for Products A and B, which is **slightly risk-averse**.

Product A	Product B
44	56

Please enter the quantities of Products A and B that you would like to order in this round.  
Total must sum to 100.

Product A	<input type="text" value="0"/>
Product B	<input type="text" value="0"/>
Total	<input type="text" value="0"/>



## RN condition

Now, in each round there will be a **decision support model** available, which recommends optimal order quantities of Products A and B. The model is based on a **complex, mathematical algorithm**.

Importantly, the algorithm is **risk-neutral** in optimizing the order mix. By that, we mean the algorithm would suggest an allocation which maximises returns **without** any consideration for risk.

**Note that whether you follow the algorithmic suggestion or not is entirely up to you.**



The **algorithm** assigned to you is

- Highly Risk-Averse
- Slightly Risk-Averse
- Risk-Neutral



**Round 9**

The table below shows the demand distribution for Products A and B in this round.

Product Type	Range of Possible Demand Values
Product A	23 - 58
Product B	8 - 133

The **algorithm** suggests the following order quantities for Products A and B, which is **risk-neutral**.

Product A	Product B
38	62

Please enter the quantities of Products A and B that you would like to order in this round.  
Total must sum to 100.

Product A

Product B

Total

**Appendix B3. Participant feedback**

Once a participant has made an order decision, a feedback screen is displayed as shown in Figure 17.



Figure 17. Feedback screen displayed to a hypothetical participant

## Appendix B4. Deviation and average profits

I report regression results for the deviation from the algorithmic recommendation as the dependent variable in Table 10. As a robustness check, I also ran regression analyses using the normalized deviation from algorithmic recommendation and the results are qualitatively similar (see Table 11). The absolute and normalized deviations were computed as follows:

$$\text{Deviation (absolute)} = \text{Subject's Allocation} - \text{Algorithmic Recommendation} \quad (21)$$

$$\text{Deviation (normalized)} = \frac{\text{Subject's Allocation} - \text{Algorithmic Recommendation}}{\text{Algorithmic Recommendation}} \quad (22)$$

**Table 10. OLS regression on (absolute) deviation from algorithmic recommendation with round-fixed effects included.**

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

	<i>Dependent variable:</i>	
	(Absolute) Deviation	
	Treatment	
	(1)	(2)
Risk-neutral	-1.365*	
	(0.744)	
Highly Risk-averse	11.889***	12.444***
	(0.921)	(0.844)
Slightly Risk-averse	3.889***	4.886***
	(0.769)	(0.676)
Age		0.051
		(0.097)
Gender		0.281
		(0.718)
Trust in Algorithm		-0.659***
		(0.222)
Previous Round Returns		0.006*
		(0.003)
Baseline Allocation		0.311***
		(0.074)
Constant		-23.451***
		(6.353)
Observations	2,760	2,622
Adjusted R <sup>2</sup>	0.605	0.359

Note: \*p<0.1; \*\*p<0.05; \*\*\*p<0.01

**Table 11. OLS regression on (normalized) deviation from algorithmic recommendation with round-fixed effects included**

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

	<i>Dependent variable:</i>	
	(Normalized) Deviation	
	Treatment	
	(1)	(2)
Risk-neutral	-1.212 (1.505)	
Highly Risk-averse	20.099*** (1.714)	19.177*** (1.708)
Slightly Risk-averse	8.929*** (1.540)	9.053*** (1.522)
Age		0.052 (0.202)
Gender		0.783 (1.480)
Trust in Algorithm		-1.640*** (0.464)
Previous Round Returns		0.011* (0.006)
Baseline Allocation		0.632*** (0.157)
Constant		-41.598*** (13.506)
Observations	2,760	2,622
Adjusted R <sup>2</sup>	0.583	0.264

*Note:* \*p<0.1; \*\*p<0.05; \*\*\*p<0.01

The intercept has been forced to 0 in column 1

Table 12 below shows subjects' average profits (per round) with standard deviations in parentheses in each of the three experimental periods.

**Table 12. Profits earned**

	<b>Baseline</b>	<b>Treatment</b>	<b>Post-treatment</b>
<b>Control</b>	744.33 (26.93)	714.03 (25.09)	645.48 (22.42)
<b>HRA</b>	746.47 (31.75)	702.66 (25.47)	642.53 (22.24)
<b>SRA</b>	740.90 (30.66)	717.32 (22.56)	645.81 (22.05)
<b>RN</b>	747.18 (29.52)	726.79 (16.07)	643.89 (25.19)

**Appendix B5. Alternate explanations**

I considered the following alternate explanations in Table 13 below, which could potentially explain the treatment and post-treatment effects that I observed in the study.

**Table 13. Alternate explanations for study results**

<b>Alternate explanation considered</b>	<b>Reason(s) for rejection</b>
Change in risk preferences	It is possible that changes in order decisions (observable) during the treatment period influenced a DM's underlying risk attitude (unobservable), which subsequently led to changes in order decisions in the post-treatment period. Such a mechanism could explain both the treatment and post-treatment effects that I observed in the study. However, I did not find any differences in risk attitude among participants assigned to different conditions in the Holt and Laury (2002) task elicited at the end of the experiment.

Loss aversion	<p>Since I use high-margin products in the study, the probability of loss in the multi-item newsvendor task is about 2%. In fact, there are no treatment rounds in which any of the algorithmic recommendations (or orders close to it) would result in losses. Hence, I do not believe that the results can be explained by loss aversion.</p>
Anchoring on mean demand heuristic	<p>I specifically considered the mean demand anchoring heuristic because it has been found to be plausible in the single-item newsvendor task. However, if subjects were to anchor on the mean demand values in the multi-item newsvendor task, their chosen order quantities would have been closer to risk-neutrality (since <math>\mu_B \cong q_{RN}</math> in the experimental setup). While this could potentially explain the ordering behavior of subjects in the control condition, the mean-demand anchoring heuristic cannot explain the treatment and post-treatment effects that I observed for subjects in the HRA algorithm condition.</p>

## Appendix C

- C1. Study 1: Algorithm design
- C2. Study 1: Task instructions
- C3. Study 1: Participant feedback
- C4. Study 1: Deviation and average profits
- C5. Study 1: PtC bias
- C6. Study 1: Demand chasing
- C7. Study 2a: Methods and Results
- C8. Study 2b: Methods and Results

### Appendix C1. Study 1: Algorithm design

Based on a meta-analysis of previous 24 newsvendor studies, the average pull-to-center (PtC) bias score was found to be 0.32 for the low-profit margin condition when the demand distribution employed was  $U(0,200)$  (See Table 2, Zhang & Siensen, 2019). The PtC bias score is computed as follows for each subject, and then averaged:

$$PtC\ score = \frac{q - q^*}{\mu - q^*} \quad (23)$$

where  $q$  is the subject's order,  $q^*$  is the profit maximizing order quantity and  $\mu$  is the average demand. A PtC score of 0.5 would mean that a subject's order is halfway between  $q^*$  and  $\mu$ . A PtC of 0.9 would mean that a subject's order quantity is closer to  $\mu$ .

An average PtC bias score of 0.32 from the meta-analysis suggests  $q = 66$ , or an over-ordering by 16 units from risk-neutral, optimal  $q^* = 50$ . Based on Feng and Gao (2020), I know that a recommendation that is 5 units less than  $q^*$  reduced over-ordering by about 4 units in their study. Thus, assuming a linear relationship, I choose a risk-averse algorithmic recommendation of  $50 - 1.25 * 16 = 30$ . To make the recommendations seem less deterministic, I then add a small noise component ( $int(U(-2,2))$ ) to generate the risk-averse algorithmic recommendation in study 1.

## Appendix C2. Study 1: Task instructions

I report below the screenshots of all task instructions from study 1 as they would appear on a participant's computer screen until round 1. I additionally report the screenshots when assigning participants to the treatment condition (RA Algorithm) as well as for the first round in the treatment period. Screenshots of the remaining parts of the study are available upon request.


**General Structure**

This study involves an inventory game that consists of 36 decision rounds.

Each round corresponds to one day in the game.

The expected time to complete the study is about 20 minutes.

**Please complete the entire study in one sitting.**



### Task Overview

Imagine you are a cart **vendor** who buys and sells fish, a highly perishable product. The cart operations process is described in the next screen.



### Cart Operations



At the beginning of each day, you will decide how much fish to order from a wholesaler to sell to customers. The unsold fish at the end of the day have to be discarded.

The unit selling price and cost of the fish are shown in the table below. They remain fixed for all the days.

Unit Selling Price (p)	Unit Cost (c)	Unit Profit (p - c)
\$12	\$9	\$3

### Customer Demand

You need to decide how much fish to order each day. However, the demand for fish is uncertain and varies every day. Importantly, **the actual demand will only be known at the end of the day** (i.e. after you place the order).

Before you order, you will be given information on the forecasted range of demand for the day. An example is shown in the table below.

Range of Forecasted Demand
0 - 200

**The following set of instructions is applicable for all days.**

**Please read them extremely carefully.**

1. The actual demand is a **random value** from the forecasted range of demand given to you for a particular day.
2. **Every integer value of demand in the range described is equally likely.** For example, based on the table above, the actual demand for fish can be any integer value in the range of 0 to 200 with equal probability.
3. The actual demand on any particular day **is unaffected** by the demand on other days.
4. You are free to order any number of fish that is within the range of forecasted demand.
5. At the end of the day, you will see your order quantity, the actual demand, cart profits, maximum possible profits and whether or not you under- or over-ordered.

## Under- and Over-Ordering



**Over-ordering:** Each day, if you end up **ordering more fish than the actual customer demand**, the excess fish will have to be discarded - i.e. you would make a **loss of \$9** for every excess unit that you ordered.

**Under-ordering:** However, if you end up **ordering less fish than the actual customer demand**, you will miss out on potential sales opportunity - i.e. you would lose out on making a **profit of \$3** for every unit of excess demand.

Therefore, for every unit you **over-order, you would lose \$9** and every unit you **under-order, you would lose out on a profit of \$3**.

**You make maximum profits when your order quantity matches the customer demand**

An example where fish is **over-ordered** is shown in the table below:  
( $p = \$12$ ,  $c = \$9$ ).

Your Order	Actual Demand	Scenario	Profits
80	60	<b>Over-ordering</b> (80 > 60)	Profits on sales: $60 \times (12 - 9) = \$180$
			Loss on excess inventory: $(80 - 60) \times 9 = \$180$
			<b>Total Profits: \$0</b> ( $180 - 180$ )

An example where fish is **under-ordered** is shown in the table below.

Your Order	Actual Demand	Scenario	Profits
60	80	<b>Under-ordering</b> (60 < 80)	Profits on sales: $60 \times (12 - 9) = \$180$
			Loss on excess inventory: \$0 (No excess)
			<b>Total Profits: \$180</b> ( $180 - 0$ )

You will be making a total of **36 decisions**, which are further divided into **3 sessions**.

### Incentives

All participants will earn Amazon vouchers worth €5 for your participation.

In addition, we will randomly select 10% of participants to win extra money based on performance. We will randomly choose 1 of the 36 decision rounds and award those lucky participants an Amazon voucher whose worth is proportional to the profits generated in that round. You can earn a maximum of €20 additionally.

**Therefore, please make your decisions carefully in every round!**

### Day 1 of 36

<b>Range of Forecasted Demand</b>
-----------------------------------

40 - 240
----------

Please enter the number of units of fish that you would like to order for this day.

units of fish



## RA Algorithm condition

### Session 2

#### Please read these instructions carefully before starting session 2

In each day of session 2, you will observe a **recommendation** regarding the optimal quantity of fish to order. The recommendation is based on a mathematical **algorithm**.



Importantly, the recommendation is **risk-averse**. That means, the algorithm tries to strike a balance between maximizing profits and generating more stable returns (i.e. lower variance in profits).

**Note that whether you follow the recommendation or not is entirely up to you.**

Day 7 of 36

Range of Forecasted Demand
39 - 239

The risk-averse algorithm suggests ordering **69** units of fish.

Please enter the number of units of fish that you would like to order for this day.

units of fish



### Appendix C3. Study 1: Participant feedback

In study 1, a risk-averse algorithmic recommendation is shown to participants assigned to the RA Algorithm condition during the treatment period (e.g., see Figure 18). Once a participant makes the order decision, a feedback screen is displayed, as shown in Figure 19.

Day 11 of 36

Range of Forecasted Demand
28 - 228

The risk-averse algorithm suggests ordering **58** units of fish.  
Please enter the number of units of fish that you would like to order for this day.

units of fish

[→](#)

Figure 18. Risk-averse algorithmic suggestion shown to participants



Figure 19. Feedback screen displayed to a hypothetical participant

## Appendix C4. Study 1: Advice utilization and average profits

**Advice Utilization.** During the treatment period, subjects in RA algorithm condition ordered about 24 more units of fish (on average) than what was recommended by the algorithm (See Figure 20 below).




Figure 20. Mean deviation from algorithmic recommendation

Weight on advice (WOA) is calculated to capture the extent of advice utilization as follows:

$$WOA = \frac{|final\ estimate - initial\ estimate|}{|advice - initial\ estimate|} \quad (24)$$

WOA typically lies between 0 and 1, with higher values indicating higher 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).<sup>25</sup> In the experiment, WOA is calculated for each subject in the RA algorithm condition in every round of the treatment period as follows:

$$WOA = \frac{q - q_c}{q_{RA} - q_c} \quad (25)$$

where  $q$  is the subject's order,  $q_{RA}$  is the RA algorithmic recommendation and  $q_c$  is the initial estimate (made without the algorithmic recommendation). Importantly,  $q_c$

---

<sup>25</sup> WOA values greater than 1 are possible, but unlikely because final estimate mostly tends to lie between the initial estimate and advice. Following prior studies, we winsorized values WOA values greater than 1 by setting each of those values equal to the 95<sup>th</sup> percentile value

represents a counterfactual scenario for what a subject in the RA algorithm condition would have ordered *before* they observed the algorithmic recommendation. I used propensity score matching to estimate  $q_c$  by pairing subjects who received the RA algorithmic recommendation (i.e., treated) with their closest counterparts who did not receive the advice (i.e., control). I used nearest neighbor propensity score matching with replacement because this technique uses all treated and control group subjects and no data is discarded by the matching process (Ho et al., 2011). The propensity score was estimated using logistic regression of the binary treatment variable (i.e., treated and control groups) on the covariates, which included average baseline order, gender, and age. This matching specification yielded good balance, as indicated by the low standardized mean differences in covariates between the matched groups (See Figure 21 below). After matching, I obtained an average WOA = 0.74.




Figure 21. Balance of covariates before and after matching in study 1

**Average profits.** Subjects in the RA algorithm condition also earned higher average profits per round (~\$72) than those in the control condition during the treatment period. However, I did not find any significant difference in profits earned by subjects between the two conditions during the post-treatment period (See Table 14 below).

**Table 14. OLS regression on profits earned with round-fixed effects included**

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

	<i>Dependent variable:</i>							
	Profits							
	Baseline		Treatment		Post-I		Post-II	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
RA Algorithm	-6.601 (26.198)	6.048 (24.321)	71.712*** (24.828)	72.298*** (22.659)	31.054 (25.945)	13.642 (25.619)	31.410 (21.688)	37.707 (25.619)
Age		0.666 (3.441)		-2.270 (3.383)		-6.108 (4.083)		0.015 (4.083)
Gender		14.072 (26.482)		12.779 (22.659)		43.741 (28.405)		15.554 (28.405)
Previous Round Profits		-0.193*** (0.032)		0.115** (0.052)		-0.231*** (0.030)		0.031 (0.030)
Baseline Order				-1.651*** (0.367)		-1.136*** (0.422)		-0.397 (0.422)
Constant	-606.560*** (46.901)	92.903 (113.939)	126.858*** (25.875)	33.821 (119.289)	-381.936*** (50.643)	449.454*** (139.993)	234.421*** (14.613)	125.618 (139.993)
Observations	456	380	1,368	1,292	456	380	456	380
Adjusted R <sup>2</sup>	0.646	0.621	0.407	0.437	0.460	0.454	0.456	0.436

Note: \*p&lt;0.1; \*\*p&lt;0.05; \*\*\*p&lt;0.01

**Appendix C5. Study 1:PtC bias**

**PtC bias.** Several previous newsvendor studies have documented that DMs tend to place orders between mean demand and optimal levels. This is called pull-to-center (PtC) bias and it can be quantified for each subject in every round as follows (See, Zhang & Siemsen, 2019).

$$PtC\ score = \frac{q - q^*}{\mu - q^*} \quad (26)$$

PtC score is a fraction typically between 0 and 1 with higher values indicating more bias (i.e., orders close to mean demand value  $\mu$ ). A PtC score of 0 indicates optimal ordering ( $q = q^*$ ). Referring to Eq. (26), I expect subjects in the RA Algorithm condition to place orders closer to  $q^*$  and exhibit lower PtC bias than those in the control condition.

Finally, since there is some debate in the literature regarding the existence of PtC bias as a phenomenon at the individual level (Lau et al., 2014b), I conducted an individual, within-subject analysis and confirmed that for a majority of subjects (~70%) in the RA Algorithm condition, the PtC score decreased from baseline to post-treatment period (See Figure 22).

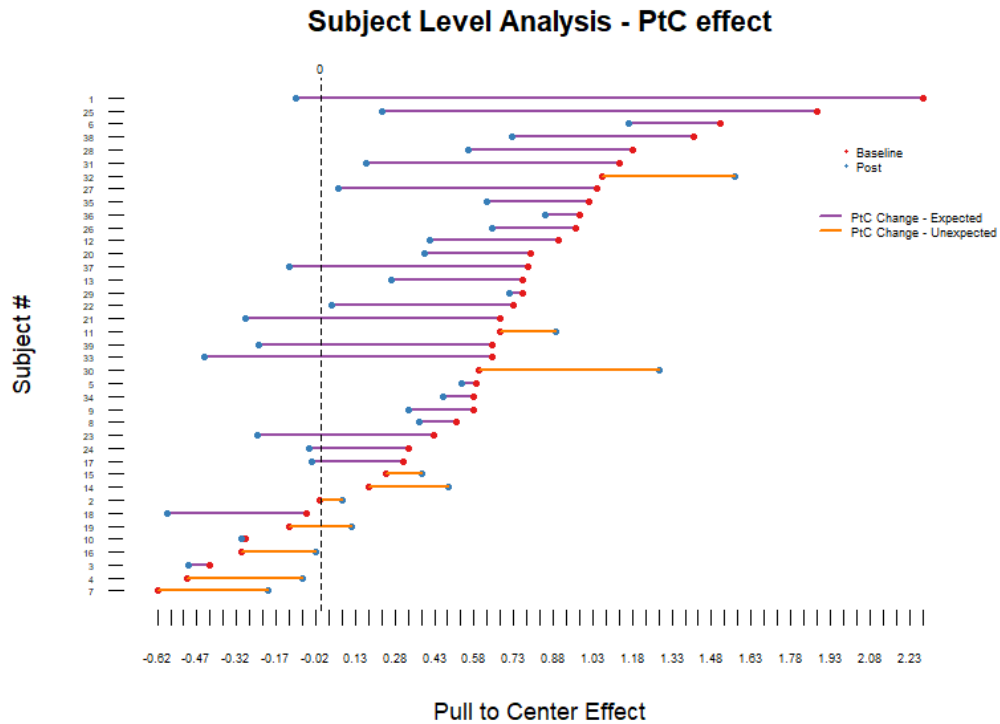


Figure 22. PtC bias - Within-subject analysis

Further, in Figure 23 below, I observe that the PtC score distribution shifts to the left (i.e., decreases) from baseline to post-treatment as a result of exposure to the risk-averse algorithmic treatment. However, there is no significant change in the PtC score distribution for subjects in the control condition.

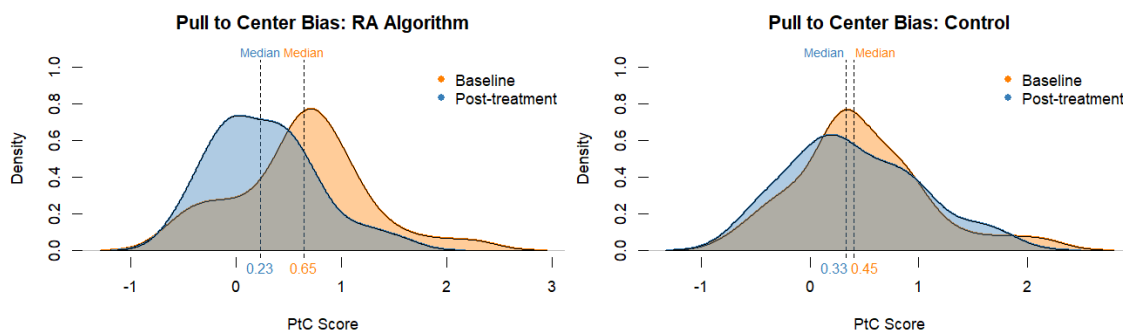


Figure 23. PtC Score Distribution

### **Appendix C6. Study 1: Demand chasing**

Following Lau and Bearden (2013), I use a simple correlation measure between  $d_{t-1}$  and  $q_t$  at the subject-level to examine demand chasing behavior in both experimental conditions during the treatment period. Under *no* true demand chasing, I expect to see a correlation score of 0 for all subjects. However, if I observe a significant positive correlation score for a number of subjects (i.e., for > 5% of subjects in either condition), I can conclude that there is evidence for demand chasing. The correlation scores between  $d_{t-1}$  and  $q_t$  at the subject-level for the treatment period is shown in Table 15 below. I do not find support for demand chasing for subjects in either experimental condition. This is expected since I varied the demand distribution in each round of the study so that the algorithmic suggestion would be different across rounds.

**Table 15. Demand chasing analysis in study 1**

subject	correlation	95% CI	condition	subjec	correlatio	95% CI	conditio
1	-0.04	[-	RA Algorithm	40	-0.18	[-0.61,0.33]	Control
2	0.20	[-	RA Algorithm	41	-0.70**	[-0.88, -	Control
3	0.09	[-	RA Algorithm	42	-0.21	[-0.63,0.3]	Control
4	0.24	[-	RA Algorithm	43	-0.06	[-0.52,0.44]	Control
5	0.51**	[0.04,0.8]	RA Algorithm	44	-0.10	[-0.55,0.4]	Control
6	0.18	[-	RA Algorithm	45	-0.24	[-0.64,0.27]	Control
7	-0.18	[-	RA Algorithm	46	-0.15	[-0.59,0.35]	Control
8	0.28	[-	RA Algorithm	47	0.17	[-0.34,0.6]	Control
9	0.07	[-	RA Algorithm	48	0.20	[-0.31,0.62]	Control
10	-0.06	[-	RA Algorithm	49	0.05	[-0.44,0.52]	Control
11	-0.26	[-	RA Algorithm	50	0.28	[-0.23,0.67]	Control
12	0.34	[-0.17,0.7]	RA Algorithm	51	0.10	[-0.4,0.56]	Control
13	-0.14	[-	RA Algorithm	52	-0.33	[-0.7,0.18]	Control
14	-0.02	[-0.5,0.47]	RA Algorithm	53	-0.15	[-0.59,0.36]	Control
15	0.51**	[0.04,0.8]	RA Algorithm	54	-0.28	[-0.67,0.23]	Control
16	-0.16	[-0.6,0.34]	RA Algorithm	55	-0.47*	[-0.77,0.02]	Control
17	0.46*	[-	RA Algorithm	56	0.25	[-0.26,0.65]	Control
18	-0.10	[-0.55,0.4]	RA Algorithm	57	-0.47*	[-0.77,0.02]	Control
19	-0.53**	[-0.81, -	RA Algorithm	58	0.18	[-0.33,0.61]	Control
20	-0.01	[-	RA Algorithm	59	0.05	[-0.44,0.52]	Control
21	0.24	[-	RA Algorithm	60	0.13	[-0.38,0.57]	Control
22	0.22	[-	RA Algorithm	61	-0.29	[-0.68,0.22]	Control
23	-0.06	[-	RA Algorithm	62	0.35	[-0.16,0.71]	Control
24	-0.16	[-0.6,0.34]	RA Algorithm	63	0.33	[-0.18,0.7]	Control
25	-0.01	[-	RA Algorithm	64	-0.06	[-0.52,0.44]	Control
26	0.28	[-	RA Algorithm	65	0.06	[-0.43,0.52]	Control
27	-0.48**	[-0.78,0]	RA Algorithm	66	0.15	[-0.36,0.59]	Control
28	0.43*	[-	RA Algorithm	67	NA	NA	Control
29	-0.10	[-0.55,0.4]	RA Algorithm	68	-0.40	[-0.74,0.1]	Control
30	0.00	[-	RA Algorithm	69	-0.07	[-0.53,0.42]	Control
31	0.00	[-	RA Algorithm	70	-0.28	[-0.67,0.23]	Control
32	-0.11	[-	RA Algorithm	71	0.19	[-0.32,0.61]	Control
33	-0.05	[-	RA Algorithm	72	-0.37	[-0.72,0.14]	Control
34	-0.72**	[-0.89, -	RA Algorithm	73	-0.04	[-0.51,0.45]	Control
35	0.11	[-	RA Algorithm	74	-0.43*	[-0.75,0.07]	Control
36	0.45*	[-	RA Algorithm	75	0.41*	[-0.08,0.75]	Control
37	0.10	[-0.4,0.56]	RA Algorithm	76	0.31	[-0.2,0.69]	Control
38	0.54**	[0.08,0.81]	RA Algorithm				
39	-0.04	[-	RA Algorithm				

Note: \*p<0.1; \*\*p<0.05; \*\*\*p<0.01;

NA values mean that the subject did not change their ordering behavior across rounds.


## Appendix C7. Study 2a: Methods and Results

### Methods

**Subjects.** As in study 1, subjects were Master's degree students. Study 2a was conducted online and subjects earned class credits in return for their participation. In addition to earning extra course credits, I selected 10% of subjects at random and awarded them Amazon vouchers based on profits earned in a randomly chosen round. The average earnings were \$12 per winning subject. The final sample size was 100 (female = 37%) with an average age of 30.5 years.

**Design and Procedure.** Study 2a was designed to closely resemble study in chapter 2 in all respects, except for one key distinction. At the end of the baseline period, subjects in study 2a were explicitly asked to choose between a highly risk-averse (HRA) algorithm and a risk-neutral (RN) algorithm as their decision support during the treatment period. The stimuli shown to participants prior to the start of the treatment period is shown below in Figure 24.

Now, in each round there will be an **algorithm** available, which recommends optimal order quantities of Products A and B. Importantly, you can choose **now** whether you want recommendations from a highly risk-averse algorithm or a risk-neutral algorithm.



The **highly risk-averse (HRA)** algorithm tries to strike a balance between maximizing profits and generating more stable returns (i.e. lower variance in profits). Therefore, it would suggest allocating **more resources** to the **less-risky product A**.

The **risk-neutral (RN)** algorithm would suggest an allocation which maximises returns without any consideration for risk.

The **profits** generated from following the suggestions of the **HRA** and **RN** algorithms will vary depending upon the actual demand values in a round.

Overall, across rounds, following the **HRA** algorithm would lead to more stable returns and following the **RN** algorithm will maximize average returns.

**Note that regardless of the algorithm you choose, whether you follow the algorithmic suggestion or not is entirely up to you.**

Which **algorithm** do you choose?

Please note that you will see recommendations from the algorithm you choose now for the next 20 rounds.

- Highly Risk-Averse
- Risk-Neutral

Figure 24. Stimuli shown to participants prior to start of treatment period in study 2a

## Results.

**Ordering behavior.** The regression results from study 2a regarding ordering behavior is shown in Table 16 below.

**Table 16. OLS regression on ordering behaviour with round-fixed effects included**

Standard errors are clustered at the subject-level.

	<i>Dependent variable:</i>					
	Allocation to Product B					
	Baseline		Treatment		Post-Treatment	
	(1)	(2)	(3)	(4)	(5)	(6)
HRA Endogenous	0.350 (1.312)	-0.130 (1.435)	-7.398*** (1.419)	-7.148*** (1.285)	-3.265** (1.321)	-3.626*** (1.235)
HRA Exogenous	-0.486 (1.258)	-0.858 (1.293)	-7.415*** (1.073)	-6.874*** (0.946)	-4.582*** (1.200)	-4.197*** (1.182)
RN Endogenous	-0.181 (1.242)	-0.125 (1.418)	-3.439*** (0.884)	-3.792*** (0.797)	-1.903* (1.116)	-1.886* (1.116)
RN Exogenous	-0.852 (1.201)	-1.030 (1.236)	-3.870*** (0.832)	-3.630*** (0.747)	-2.513** (1.184)	-1.979* (1.130)
Age		0.044 (0.108)		0.097 (0.085)		0.112 (0.094)
Gender		-1.105 (0.787)		0.388 (0.621)		0.493 (0.683)
Previous Round Profits		0.015*** (0.005)		0.006** (0.003)		0.003 (0.004)
Baseline Allocation				0.380*** (0.065)		0.403*** (0.060)
Constant	65.819*** (1.045)	54.708*** (5.337)	64.115*** (0.955)	35.225*** (5.831)	70.936*** (1.012)	33.473*** (5.219)
Observations	1,896	1,659	4,740	4,503	2,844	2,607
Adjusted R <sup>2</sup>	0.117	0.127	0.138	0.169	0.070	0.109

Note: \*p<0.1; \*\*p<0.05; \*\*\*p<0.01

**Advice Utilization.** I computed weight of advice (WOA) score for each treated subject to capture the extent of advice utilization. Weight on advice (WOA) is calculated to capture the extent of advice utilization as follows:

$$WOA = \frac{|final\ estimate - initial\ estimate|}{|advice - initial\ estimate|} \quad (27)$$

WOA typically lies between 0 and 1, with higher values indicating higher 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).<sup>26</sup> In the experiment, WOA is calculated for each subject in the treated condition in every round of the treatment period as follows:

$$WOA = \frac{q - q_c}{q_T - q_c} \quad (28)$$

where  $q$  is the subject's order,  $q_T$  is the algorithmic recommendation and  $q_c$  is the initial estimate (made without the algorithmic recommendation). Importantly,  $q_c$  represents a counterfactual scenario for what a subject in a treated condition would have ordered *before* they observed the algorithmic recommendation. I used propensity score matching to estimate  $q_c$  by pairing subjects who received an algorithmic recommendation (i.e., treated) with their closest counterparts who did not receive the advice (i.e., control). I used nearest neighbor propensity score matching with replacement because this technique uses all treated and control group subjects and no data is discarded by the matching process (Ho et al., 2011). The propensity score was estimated using logistic regression of the binary treatment variable (i.e., treated and control groups) on the covariates, which included average baseline order, gender, and age. This matching specification yielded good balance, as indicated by the low standardized mean differences in covariates between the matched groups (See Figure 25 below).

---

<sup>26</sup> WOA values greater than 1 are possible, but unlikely because final estimate mostly tends to lie between the initial estimate and advice. Following prior studies, we winsorized values WOA values greater than 1 by setting each of those values equal to the 95<sup>th</sup> percentile value.

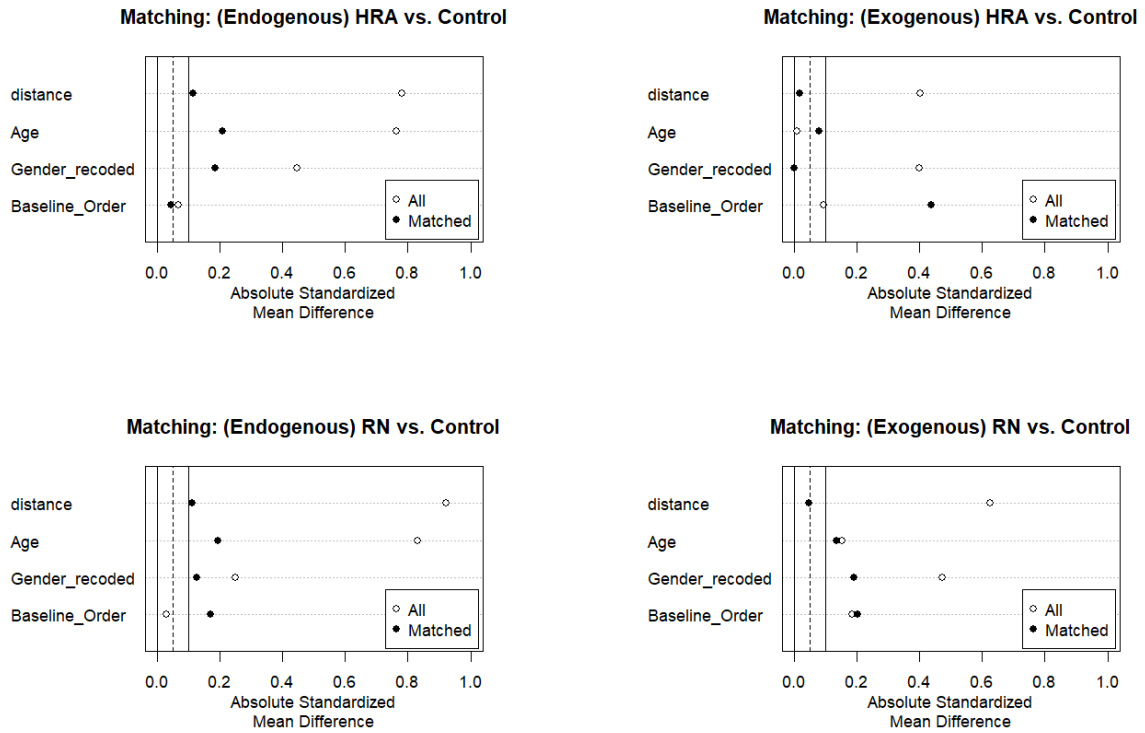


Figure 25. Balance of covariates before and after matching

Regardless of whether the (HRA or RN) algorithm is assigned externally (i.e., exogenous assignment) or chosen by the subjects themselves (i.e., endogenous assignment), I find similar levels of advice utilization. Additionally, I find that subjects who encountered the RN algorithm utilized the advice about 30% more than those who observed the HRA algorithm. See Figure 26 and Table 17.

*Note that each blue dot represents the average WOA score computed for a subject during the treatment period*




Figure 26. WOA Distribution

**Table 17. OLS regression on weight of advice with round-fixed effects included.** Standard errors are clustered at the subject-level.

	<i>Dependent variable:</i>
	Weight of Advice WOA
Exogenous	-0.006 (0.023)
RN Algorithm	0.299*** (0.020)
Baseline Allocation	-0.002 (0.002)
Age	-0.008*** (0.003)
Gender	0.014 (0.024)
Previous Round Profits	-0.0004*** (0.0001)
Constant	1.260*** (0.182)
Observations	3,548
Adjusted R <sup>2</sup>	0.185

*Note:* \*p<0.1; \*\*p<0.05; \*\*\*p<0.01

## Appendix C8. Study 2b: Methods and Results

### Methods

**Subjects.** Subjects were MBA students from a large private university. Study 2b was conducted online and I used a similar incentive scheme employed in previous study 1. The final sample size was 161 (female = 38 %) with an average age of 29.7 years. The average earnings were \$24 per winning subject.

**Design and Procedure.** Study 2b was designed to be similar to the study in chapter 2 in all important respects except for one key distinction<sup>27</sup>. At the end of the baseline period, subjects were randomly assigned to one of two treatment conditions in which they observed recommendations that were generated by a highly risk-averse algorithm

<sup>27</sup> Study 2b had 36 rounds, which were designed to be identical to 36 rounds in chapter 2 study.

(*HRA Algorithm*) or a highly risk-averse human participant (*HRA Human*) who had taken part in a similar study previously. The study design is shown in Figure 27 below.

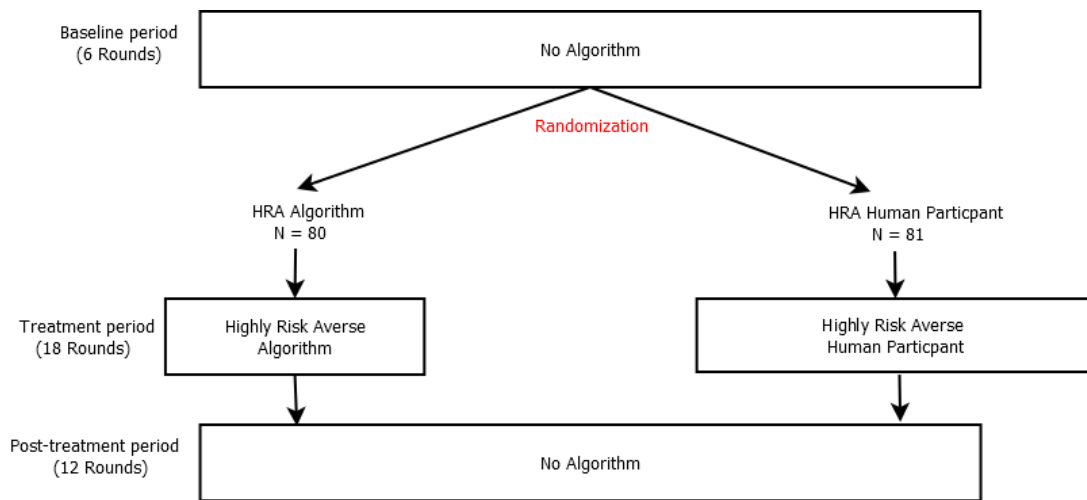


Figure 27. Study 2b design

**Human and algorithmic recommendations.** The objective was to design HRA human and algorithmic recommendations that were equivalent in terms of level of risk aversion. To achieve that, I first chose subject(s) in the control condition in each treatment round from the study in chapter 2 whose order decisions most closely matched those of the HRA Algorithm that I employed in the study. The order decisions ( $q_{Bt}$ ) made by those subject(s) served as the HRA Human recommendations for study 2b. Next, I back-calculated the level of risk aversion ( $\lambda_t$ ) in each round to which these order decisions ( $q_{Bt}$ ) correspond, by assuming that the HRA Algorithm maximizes an objective function of the form  $E(\pi) - \lambda V(\pi)$ , as previously discussed. Then, I found the average level of risk aversion ( $\lambda$ ) and used this as the risk aversion parameter for the HRA Algorithm. I therefore ensured that the HRA Human and HRA Algorithm recommendations were identical in terms of average risk aversion level.

Subjects in both the treatment conditions were free whether or not to follow the algorithm (or human participant) recommendation. The stimuli shown to participants prior to random assignment to either of the treatment conditions is shown below in Figure 28.

## Session 2

### Please read these instructions carefully before starting session 2

In each round of session 2, you will observe a **recommendation** regarding the optimal order quantities of Products A and B. The recommendation is based either on a mathematical **algorithm** or on the response of another (human) **participant** from an identical study conducted previously among MBA students.

Importantly, the recommendations of both algorithm and human participant are **equally highly risk-averse**. That means the recommendations typically suggest to **allocate more resources** to the **less-risky product A**.

On the next screen you will be informed whether you have been assigned to receive the **algorithmic** or **human participant** recommendation.

**Note that whether you follow the recommendation or not is entirely up to you.**

Figure 28. Sample stimuli used in study 2b

## **Results**

**Ordering behavior.** The regression results from study 2b regarding ordering behavior is shown in Table 18 below. Additionally, I find that subjects are twice as more likely to prefer an algorithm than a human recommendation if given a chance to repeat the task (See Figure 29).

**Table 18. OLS regression ordering behavior with round-fixed effects included**

Standard errors are clustered at the subject-level

	<i>Dependent variable:</i>					
	Allocation to Riskier Product B					
	Baseline		Treatment		Post-Treatment	
	(1)	(2)	(3)	(4)	(5)	(6)
HRA Human	0.781 (0.910)	0.760 (0.981)	1.830** (0.888)	1.675** (0.791)	0.294 (1.035)	-0.137 (0.994)
Control <sup>28</sup>	0.487 (1.103)	1.253 (1.286)	6.153*** (0.963)		3.084*** (1.171)	2.620** (1.169)
Age		0.138 (0.165)		0.039 (0.158)		-0.030 (0.153)
Gender		1.241 (0.901)		0.574 (0.804)		0.501 (0.938)
Previous Round Profits		0.008 (0.005)		0.013*** (0.003)		0.009** (0.004)
Baseline Allocation				0.339*** (0.069)		0.341*** (0.073)
Trust in Recommendation				-0.732** (0.328)		
Constant	64.567*** (0.724)	55.105*** (5.851)	59.583*** (0.824)	30.346*** (7.142)	68.441*** (0.883)	39.997*** (6.568)
Observations	1,242	1,035	3,726	2,737	2,484	2,277
Adjusted R <sup>2</sup>	0.145	0.156	0.110	0.132	0.057	0.087

*Note:* \*p<0.1; \*\*p<0.05; \*\*\*p<0.01<sup>28</sup> The control condition is from the study in chapter 2.

Note: Error bars signifies  $\pm 1$  standard deviation




Figure 29. Recommendation preference

**Advice Utilization.** As in previous study, I computed weight of advice (WOA) score for each treated subject using nearest neighbor propensity score matching with average baseline order, age, and gender as covariates for the matching process. The matching specification yielded good balance, as indicated by the lowered mean distance scores (See Figure 30 below). I did not find any difference in advice utilization between HRA Human and HRA Algorithm conditions (See Figure 31 and Table 19).




Figure 30. Balance of covariates before and after matching in study 2b

Note that each blue dot represents the average WOA score computed for a subject during the treatment period.




Figure 31. WOA Distribution

**Table 19. OLS regression on weight of advice with round-fixed effects included.**  
Standard errors are clustered at the subject-level.

	<i>Dependent variable:</i>
	Weight of Advice WOA
HRA Human	-0.010 (0.022)
Baseline Allocation	-0.005*** (0.002)
Age	-0.004 (0.004)
Gender	-0.046** (0.023)
Previous Round Profits	-0.0004*** (0.0001)
Constant	1.373*** (0.177)
Observations	2,681
Adjusted R <sup>2</sup>	0.040
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01

## References

- Abdellaoui, M., Diecidue, E., & Öncüler, A. (2011). Risk preferences at different time periods: An experimental investigation. *Management Science*, 57(5), 975–987. <https://doi.org/10.1287/mnsc.1110.1324>
- 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>
- Bastani, H., Bastani, O., & Sinchaisri, P. (2021). Learning Best Practices: Can Machine Learning Improve Human Decision-Making? *Academy of Management Proceedings*, 2021(1), 14006. <https://doi.org/10.5465/ambpp.2021.14006abstract>
- 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>
- Bertsimas, D., & Kallus, N. (2016). *The Power and Limits of Predictive Approaches to Observational-Data-Driven Optimization*.

- Bhatia, S. (2012). *The Dynamics of Anchoring in Bidirectional Associative Memory Networks*.
- Bhattacharya, U., Hackethal, A., Kaesler, S., Loos, B., & Meyer, S. (2012). Is unbiased financial advice to retail investors sufficient? Answers from a large field study. *Review of Financial Studies*, 25(4), 975–1032.  
<https://doi.org/10.1093/rfs/hhr127>
- Blavatsky, P., Ortmann, A., & Panchenko, V. (2021). On the Experimental Robustness of the Allais Paradox. *American Economic Journal: Microeconomics*. <https://doi.org/10.1257/MIC.20190153>
- 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.  
<https://doi.org/10.1287/msom.1060.0190>
- Bolton, G. E., Ockenfels, A., & Thonemann, U. W. (2012). Managers and Students as Newsvendors. *Management Science*, 58(12), 2225–2233.  
<https://doi.org/10.1287/mnsc.1120.1550>
- Caro, F., Colliard, J.-E., & Katok, E. (2021.). *Call for Papers-Management Science Special Issue on the Human-Algorithm Connection*  
<https://doi.org/10.1287/mnsc.2021.4260>
- 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>
- Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-Dependent Algorithm Aversion. *Journal of Marketing Research*, 56(5), 809–825.  
<https://doi.org/10.1177/0022243719851788>

- 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).  
<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>
- Davis, A. M., Flicker, B., Hyndman, K., Katok, E., Keppler, S., Leider, S., Long, X., & Tong, J. D. (2023). A Replication Study of Operations Management Experiments in Management Science. *Management Science*, 69(9), 4977.  
<https://doi.org/10.1287/mnsc.2023.4866>
- 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>
- Eeckhoudt, L., Gollier, C., & Schlesinger, H. (1995). The Risk-Averse (and Prudent) Newsboy. *Management Science*, 41(5), 786–794.  
<https://doi.org/10.1287/mnsc.41.5.786>
- 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*.  
<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>
- Filiz, I., Judek, J. R., Lorenz, M., & Spiwoks, M. (2021). Reducing algorithm aversion through experience. *Journal of Behavioral and Experimental Finance*, 31, 100524. <https://doi.org/10.1016/j.jbef.2021.100524>

- Föbus, N., Lange, T., Leopoldseder, M., & Magnus, K.-H. (2019). *Food retail: The path to autonomous planning | McKinsey*.
- Fox, C. R., & Tversky, A. (1995). Ambiguity aversion and comparative ignorance. *Quarterly Journal of Economics*, 110(3), 585–603.  
<https://doi.org/10.2307/2946693>
- Franco, L. A., Hämäläinen, R. P., Rouwette, E. A. J. A., & Leppänen, I. (2021). Taking stock of behavioural OR: A review of behavioural studies with an intervention focus. *European Journal of Operational Research*, 293(2), 401–418.  
<https://doi.org/10.1016/J.EJOR.2020.11.031>
- 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>
- 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>

- Kirshner, S. N., & Moritz, B. B. (2022). For the future and from afar: Psychological distance and inventory decision-making. *Production and Operations Management*. <https://doi.org/10.1111/POMS.13829>
- Koehler, D. J., & Beaugard, 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>
- Lau, N., Hasija, S., & Bearden, J. N. (2014). Newsvendor pull-to-center reconsidered. *Decision Support Systems*, 58(1), 68–73. <https://doi.org/10.1016/J.DSS.2012.12.041>
- 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>

- Luo, X., Tong, S., Fang, Z., & Qu, Z. (2019). Frontiers: Machines vs. humans: The impact of artificial intelligence chatbot disclosure on customer purchases. *Marketing Science*, 38(6), 937–947. <https://doi.org/10.1287/mksc.2019.1192>
- Lurie, P. M., & Goldberg, M. S. (1998). An approximate method for sampling correlated random variables from partially-specified distributions. *Management Science*, 44(2), 203–218. <https://doi.org/10.1287/mnsc.44.2.203>
- Mahmud, H., Islam, A. K. M. N., Ahmed, S. I., & Smolander, K. (2022). What influences algorithmic decision-making? A systematic literature review on algorithm aversion. *Technological Forecasting and Social Change*, 175, 121390. <https://doi.org/10.1016/j.techfore.2021.121390>
- Markowitz, H. (1952). Portfolio Selection\*. *The Journal of Finance*, 7(1), 77–91. <https://doi.org/10.1111/J.1540-6261.1952.TB01525.X>
- 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>
- Mehra, R., & Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2), 145–161. [https://doi.org/10.1016/0304-3932\(85\)90061-3](https://doi.org/10.1016/0304-3932(85)90061-3)
- Newman, D. T., Fast, N. J., & Harmon, D. J. (2020). When eliminating bias isn't fair: Algorithmic reductionism and procedural justice in human resource decisions. *Organizational Behavior and Human Decision Processes*, 160, 149–167. <https://doi.org/10.1016/j.obhdp.2020.03.008>
- Prakash Katariya, A., Cetinkaya, S., Tekin, E., & Cetinkaya, 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>

Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers and Artificial Intelligence: An Experiential Perspective. *Journal of Marketing*, 85(1), 131–151.

<https://doi.org/10.1177/0022242920953847>

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>

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>

Snizek, 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>

Thaler, R. H., & Johnson, E. J. (1990). Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice. *Management Science*, 36(6), 643–660. <https://doi.org/10.1287/mnsc.36.6.643>

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*, 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., Choshen-Hillel, S., & Milyavsky, M. (2009). *Spurious Consensus and Opinion Revision: Why Might People Be More Confident in Their Less Accurate Judgments?* *Journal of Experimental Psychology: Learning Memory and Cognition*.

<https://doi.org/10.1037/a0014589>

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>