ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE
DELLA RICERCA

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Free-riding product returns to drive profits

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:
Montaguti, E., Neslin, S.A., Valentini, S. (2025). Free-riding product returns to drive profits. JOURNAL OF MARKETING, -, 1-50 [10.1177/00222429251409673].

Availability:
This version is available at: https://hdl.handle.net/11585/1042317 since: 2026-02-05

Published:
DOI: http://doi.org/10.1177/00222429251409673

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (https://cris.unibo.it/).
When citing, please refer to the published version.

(Article begins on next page)

05 June 2026

JOURNAL of Marketing

Free-Riding Product Returns to Drive Profits

Journal:	*Journal of Marketing*
Manuscript ID	JM-23-0779.R4
Manuscript Type:	Revised Submission
Research Topics:	Customer Relationship Management, Relationship Marketing, Sales Promotion
Methods:	Logit/Probit/Tobit, Optimization, Panel Data Models, Regression Models, Secondary Data Analysis, Lab Experiments

SCHOLARONE™ Manuscripts

Title:

Free-Riding Product Returns to Drive Profits

Authors:

Elisa Montaguti
Full Professor of Marketing, Management Department, University of Bologna, Via Capo di Lucca, 34 40126 Bologna, Italy, Phone (+39) 051 209 8089, Fax: (+39) 051 246 411, email: elisa.montaguti@unibo.it

Scott Neslin
Albert Wesley Frey Professor of Marketing, Tuck School of Business, Dartmouth College, Hanover, NH 03755, USA, Phone: (+1) (603) 646 2841, Fax: (+1) (603) 646 1308, email: scott.a.neslin@dartmouth.edu

Sara Valentini
Associate Professor, Marketing Department, Bocconi University, Italy, Via Röntgen, 1, 20136 Milano, Italy Phone: (+39) 02 58362487 Email: s.valentini@unibocconi.it

Peer Review Version

Data Availability Statement

An alternative disclosure plan was approved for this article.
The data that support the findings of this article are not available due to a nondisclosure agreement with the partnering firm. However, replication materials (including a synthetic dataset that mimics the structure of the original data and all analysis code) have been uploaded to the Journal of Marketing Dataverse in restricted form and are available from the authors on reasonable request, subject to the terms of the NDA.

Acknowledgment

The authors thank Asim Ansari (Columbia University) for his insightful suggestions on issues related to the identification of dynamic coefficients, Jim Smith (Tuck School of Business) for his invaluable guidance in implementing the optimization, and Joachim Vosgerau (Bocconi University) for his helpful comments during the development of this project. We also thank participants at the Marketing Science and EMAC conferences for their insightful suggestions and comments.
.

Free-Riding Product Returns to Drive Profits

Abstract

Returning products has become standard practice for consumers, and a significant "pain point" for retailers. The authors contend that returns can be harnessed to increase profits. This requires retailers to manage the interweaving dynamics of product returns and purchase. The strategy is to co-manage a virtuous cycle whereby current returns increase future purchases, and a vicious cycle whereby current returns increase future returns. Marketing executes the strategy. It generates purchases. The returns that follow impose a direct cost due to the vicious cycle. However, the virtuous cycle of returns can offset this, enabling the retailer to "free-ride" returns by optimizing its marketing. The approach follows the Decision Support System (DSS) paradigm by combining a conceptual model, a statistical model, data, and optimization. A core construct is a stock variable tracking consumers' memory of return experiences. This drives both the virtuous and vicious cycles. The authors optimize marketing spend accounting for Return Stock. The best results are when dynamics are included in both the statistical and optimization models. Results suggests managers should avoid strict return policies aimed at eliminating returns. Instead, they should design policies that optimally balance the long-term benefits and costs of returns.

Keywords: product returns, returns carryover effect, dynamic optimization, marketing ROI, promotions, Return Stock, return policies, online retailing, marketing spend allocation

Peer Review Version

Product returns are the bane of retail management. Retailers bemoan the deduction from revenues they thought were in the bank and the costs of processing returned merchandise. The problem is acute in online retailing, where return rates easily approach 30-40%.

There is no doubt that returns saddle retailers with a short-term reduction in profits. However, research suggests returns may benefit retailers in the long run by increasing future sales. Bower and Maxham (2012) find that customers who make returns increase spending in the long run if return shipping costs are free. In an extensive empirical analysis, Petersen and Kumar (2015) find a positive relationship between current returns and future sales. Chen (2021) develops a structural model to investigate the returns/sales relationship and replicates the positive relationship. El Kihal et al. (2025) uncover this relationship in their research. In a meta-analysis, Janakiraman et al. (2016), find lenient return policies unsurprisingly increase returns, but increase purchase to a greater extent. While the psychological mechanism behind this virtuous cycle has not been definitively identified and boundary conditions are not ironed out, the evidence for a positive relationship between current returns and future purchases is compelling.

This evidence is provocative because it suggests retailers have an opportunity to manage returns to maximize long-term profits. How can this be achieved? Our work suggests firms can improve long-term profits by devising a dynamic marketing strategy that accounts for the relationship between product returns and future purchases. We refer to this as "free-riding" product returns. Free-riding requires optimizing marketing over time, dynamically.

Complexity arises due to dynamics among marketing, returns, and purchase. Marketing can produce a purchase in period *t* that may get returned in period *t*. Due to the positive impact

Peer Review Version

of returns on purchase, this would increase the chance of a purchase in t+1.[1] This is the *virtuous* cycle. However, it could make a return more likely in t+1. This is the *vicious* cycle. If the virtuous cycle is large relative to the vicious cycle, we can free-ride the period t return and let it generate a purchase in period t+1. We don't have to market in t+1. We can market in t+2 to generate a purchase and go through the sequence again.

Marketing plays a pivotal role as it can stimulate purchases that can be returned. These returns can beget future purchases due to the virtuous cycle as well as more returns due to the vicious cycle. Therefore, marketing stimulates both the vicious and virtuous cycles. The "trick" is to use marketing to generate enough returns to free-ride the virtuous cycle, but not too many returns, which would amplify the vicious cycle. This requires a statistical model that describes these dynamics, and a dynamic optimization that increases long-term profits.

The purpose of this paper is to show how retailers can use marketing to optimize long-term profits by managing returns. Our goal is to: i) develop a framework that describes the relationships among current returns, future returns and purchases, the interplay between these vicious and virtuous cycles, and the role of marketing in these dynamics; ii) develop and estimate the statistical model that follows from this framework; iii) with the estimates as inputs, use dynamic programming to generate an optimal free-riding policy; iv) gain insights into how dynamic optimization increases profits. Is it by free-riding the future benefits of returns? How does this affect marketing spend, scheduling and allocation across activities? How does the free-riding strategy interact with the firm's product return policy?

By combining model, data, statistics, and optimization, our approach follows the decision support system paradigm (DSS) (Little 1979; Lilien, Kotler, and Moorthy 1992). The

[1] We sometimes use words like "influence", "impact", and "affect" that suggest causal relationships. We do this to be concise. Our results should be interpreted as associative and correlational. They do not prove causality.

Peer Review Version

model is the framework.[2] It identifies key phenomena – virtuous and vicious cycles – and how they relate to marketing. The *data* are customer-level purchases and returns. The *statistics* are equations that mathematically represent the framework and the econometrics required to estimate them. The estimates serve as the descriptive tool at the core of the DSS. The *optimization* is a stochastic dynamic program. A DSS produces tangible recommendations and insights that guide decision-making. For example, in our study, we recommend concrete marketing scheduling and allocation. One insight is that the goal is *not* zero returns – it is to use marketing to generate enough returns to free-ride the virtuous cycle without generating too much negative impact – the vicious cycle.

Our data cover a five-year period and include approximately 1,500 customers from several countries, observed from their first purchase and known to have made at least one purchase per year. The data are provided by an online fashion retailer that has a "standard" 30% return rate. We show the role customers' memories of return experience play in virtuous and vicious cycles and measure the impact of marketing on these cycles. We examine how optimal marketing varies across countries, which, due to different return policies, differ in their potential to generate profits through free-riding.

The estimated statistical model finds a positive direct impact of current returns on future sales, the virtuous cycle. It also finds a positive direct impact of current returns on future returns, the vicious cycle. We show that dynamic optimization can prescribe a marketing policy that pulses spending to exploit the virtuous cycle without over-generating the vicious cycle.

We show that dynamic optimization outperforms myopic optimization that does not account for the future benefit of returns, and the benefit of dynamic optimization is not realized

[2] The term "model" here is conceptual in the sense of Little (p. 10) who defines it as "a preconceived idea of how the world works". We will use the term "statistical model" to refer to the equations that express this mathematically. In sum, we translate the conceptual model into a statistical model we estimate with our data.

Peer Review Version

unless the statistical model also incorporates dynamics. We show that a successful free-riding strategy may skew marketing spend toward activities that have a large impact on sales but a smaller direct impact on returns. For example, free (outbound) shipping is problematic because it has a substantial impact both on sales and returns. The dynamic optimization therefore re-allocates marketing spend away from free shipping and toward advertising that builds sales without a large direct effect on returns.

We show the firm's return policy influences the profitability of free-riding. Strict policies reduce returns, but in so doing also reduce the opportunity to free-ride returns. Lenient policies can strike a better balance between the long-term benefits and costs of returns, opening the door to free-riding.[3] For instance, when our focal company implements strict return policies (e.g., in Denmark, Greece, and Russia), the difference between the benefits and costs of returns – between virtuous and vicious cycles – narrows, limiting the opportunity to capitalize on returns. In contrast, more lenient policies, as seen in the USA, increase this difference, offering greater opportunity to free-ride the way to higher profits.

This paper contributes in three ways: *integrating* existing knowledge, generating *insights*, and showing how to *implement* the free-riding strategy.

First, we *integrate* a statistical model of virtuous and vicious cycles with a dynamic optimization to produce profitable free-riding. To our knowledge, no previous research has done this.

Second, in line with the DSS approach, we offer *insights* into how to manage free-riding. We show how and why lenient return policies are associated with higher levels of free-riding. We find free-riding works not necessarily by reducing marketing spend but by carefully

[3] We thank the anonymous reviewers for their insightful suggestion to examine regional differences, allowing us to distinguish regions with lenient versus strict return policies.

Peer Review Version

scheduling and allocating it among multiple marketing activities. We show the evolution of consumers' memories of past return experiences, quantified by "Return Stock", is the fulcrum for generating virtuous and vicious cycles; it therefore enables successful free-riding. We suggest the mechanisms for translating Return Stock into virtuous and vicious cycles are by enhancing consumer perceptions of fairness and their relationship with the company, while reducing consumers' concerns of purchase risk (buying the wrong product).

Also in line with the DSS paradigm, we investigate the key requirement for *implementing* the free-riding strategy. It is to include return dynamics in *both* the statistical and optimization models. We show our formulation of Return Stock and our dynamic stochastic optimization generates the largest free-riding profits. We find that successful implementation requires a balance between virtuous and vicious cycles. Implementation should produce a profitable level of returns, not zero returns.

In summary, we show how a novel dynamic statistical model optimized over time can be applied in real-world contexts to free-ride the long-term benefits of product returns. By combining model, statistics, data, and optimization, our approach functions as a DSS. We show free-riding requires dynamics both in the statistical model and in optimization. A dynamic optimization paired with a non-dynamic statistical model lacks the necessary long-term effects; a dynamic statistical model with myopic optimization fails to fully leverage long-term effects. A fully dynamic approach provides a practical DSS for optimizing return management. By striking a balance that avoids overly producing the vicious cycle while still exploiting the benefits of returns, firms can achieve both financial and environmental gains.

Literature Review

Peer Review Version

We summarize research in marketing focusing on three streams relevant to our work: (1) the impact of returns on future purchase (the virtuous cycle), (2) the impact of returns on future returns (the vicious cycle), and (3) the impact of marketing on returns.

The virtuous cycle: Emerging evidence suggests returns can increase future purchases and revenues (Petersen and Kumar, 2009, 2015; Bower and Maxham 2012; Chen 2021). This is intriguing and apparently not well accounted for by practitioners as well as academics. Previous research suggests two possible explanations for this virtuous cycle: risk reduction (Petersen and Kumar 2015; Chen 2021) and fairness (Bower and Maxham 2012).

Many times, consumers risk dissatisfaction by buying the "wrong" product. Anderson, Hansen and Simester (2009) conceptualize returns as the consumers' economic option consumers to remedy this risk. This makes clear that the ability to return encourages a customer to purchase *now*. But why should the return option increase purchasing in the future?

One possibility is that product returns reduce consumers' *future* risk. Petersen and Kumar (2009) posit that returns enable consumers to learn more about the retailer's products, thus reducing the risk of buying the wrong product in the future. Venkatesan and Kumar (2004) posit that returns provide firms with the opportunity to increase customer satisfaction. However, too many returns may prevent the firm from taking advantage of this opportunity. Venkatesan and Kumar (2004) as well as Petersen and Kumar (2009) find an inverted-U relationship between returns and future purchases.

Bower and Maxham (2012) show that if customers perceive the retailer handles returns fairly, returning product produces less regret and consumers even experience "rejoicing or elation" (p. 133). This increases customer satisfaction and generates more purchases in the future. If consumers perceive returns to be unfairly managed, they regret the purchase, presumably making them dissatisfied and decreasing future purchase.

Peer Review Version

In sum, research to date suggests that risk reduction and perceived fairness play a role in the empirical finding that returns beget future purchases. More work is needed to probe these explanations as well as others. The key takeaway, however, is that returns can exert a positive impact on future purchases.

Petersen and Kumar (2015) leverage this finding to develop a CLV model accounting for product returns. They estimate this model and optimize marketing spend one period at a time. This "myopic optimization" indicates that retailers can increase profits by accounting for benefits and costs of returns. Their empirical model includes dynamics, but they do not optimize dynamically. This we propose is a key desideratum. The optimal marketing strategy might tolerate an acceptable level of returns in the short run to free-ride the virtuous cycle. This requires the optimization to look ahead to incorporate the impact of current actions on future profits. One needs dynamic empirical *and* optimization models.

The vicious cycle: El Kihal, Erdem, Schulze and Zhang (2025) investigate how return rates evolve over time. Returns provide customers with an opportunity to learn about the brand's product quality and fit, decreasing future returns. Return rates, however, could increase because customers develop a return habit or become more discerning over time. El Kihal et al. (2025) find return rates *increase* over time. This suggests the "vicious cycle" whereby current returns beget future returns. This is important as it demonstrates the direct impact of current returns on future returns must be included in any empirical model.

The impact of marketing on returns: Shehu, Papies, and Neslin (2020) suggest free shipping promotions are a risk premium that encourages risky purchase and a commensurate increase in returns. El Kihal and Shehu (2022) posit marketing elevates both the expected and experienced value of a purchased product. If the expected value is less than the experienced value, the product is likely to be returned. The authors found several marketing instruments

Peer Review Version

increased returns. This suggests that marketing should be included in modeling the interplay among marketing, purchases, and returns.

Table 1 summarizes the above discussion and shows what distinguishes this paper. Our work is unique in integrating a dynamic descriptive model and a dynamic optimization of marketing. To our knowledge, no one else has done this. It is essential because as we will show, not considering one of these dynamics short-changes the other. In addition, our work is uniquely comprehensive. Specifically, it (1) Investigates mediators that serve as mechanisms for translating returns into future behavior, (2) Considers the impact of multiple marketing activities on returns, (3) Conceptualizes "Return Stock" to quantify consumer memory of the returns experience, (4) Empirically studies the impact of different return policies. While some previous papers have investigated a subset of these issues, none has investigated all four of them.

Our work can thus be viewed as a Decision Support System (DSS) (Little 1979). It combines models, statistics, data, and optimization to generate tangible marketing plans and conceptual insights for free-riding returns. For example, we show the gap between the magnitudes of virtuous and vicious cycles provides room for free-riding – strict return policies minimize this gap; lenient policies increase it. Therefore, lenient policies provide the best opportunity for free-riding. Capturing these dynamics is important, yet to our knowledge, no prior work has investigated this. Our approach explicitly models these inherently dynamic cycles and optimizes decisions accordingly.

For Peer Review Version

Table 1 Literature Review Summary

Research	Returns → Future Purchase (*Virtuous Cycle*)	Returns → Future Returns (*Vicious Cycle*)	Mediating the Virtuous and Vicious Cycles	Number of Marketing Activities Influencing Returns	Quantifying the Returns Experience	Return Policy	Dynamic Descriptive Model	Dynamic Optimization of Marketing	Dynamic Model + Dynamic Optimization
Petersen and Kumar 2009	√	√		Single	Total Returns	Single	Yes	Myopic	No
Petersen and Kumar 2015	√	√		Single	One-period Lag	Single	Yes	Myopic	No
Chen 2021	√	√		Single		Single	Yes		
El Kihal et al. 2025	√	√		Multiple	Stock Variable		Yes		
El Kihal and Shehu 2022				Multiple					
Shehu et al 2020				Single			Yes		
Wood 2001	√	√				Strict vs. Lenient			
Bower and Maxham 2012	√		Fairness			Free vs. Fee			
This Paper	√	√	**Fairness, Relationship, Risk**	**Multiple**	**Stock Variable**	**Strict vs. Lenient**	**Yes +**	**Yes →**	**Free-Riding**

Empty cells indicate aspects not addressed by the respective study.

Peer Review Version

Conceptual Framework

The above discussion suggests basis for our framework, portrayed in Figure 1. The framework depicts the relationships among returns, marketing, and customer-level sales. Figure 1 shows that marketing induces both sales and returns, as documented by our literature review. The role of returns is captured by "Return Stock." The Engel, Blackwell, and Miniard (EBM) model of consumer decision-making (1995, p. 152) motivates this concept. Return Stock is a key component of our framework; hence we elaborate.

Figure 1: Conceptual Model - Framework.

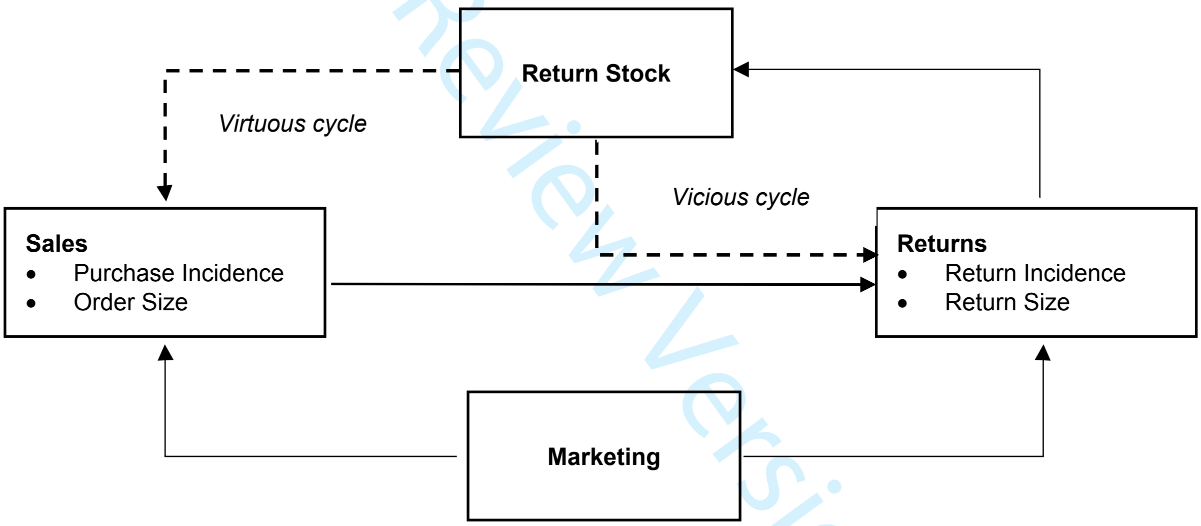

EBM posits that memory governs beliefs, attitudes, and intentions regarding future behavior. Return Stock captures that memory. We are concerned with two behaviors – returns and purchase. When considering whether to return, customers assess their beliefs about factors such as ease of returning and hassle costs (Bower and Maxham 2012). If, based on memory, these beliefs are favorable, they foster attitude and intention that encourage returns, creating the vicious cycle. However, memory-supported beliefs that the retailer handles returns fairly and expeditiously create favorable attitudes toward the retailer, increasing future purchases. This is the virtuous cycle.

Peer Review Version

The following factors help create memory: First, memory is strengthened by multiple experiences. Second, the more recent these experiences are, the stronger the memory. Third, a memory not strengthened by multiple recent experiences fades away. For instance, if a customer returns four items within three months, they will have a strong memory of returning products. Conversely, if it's been three years since the last return, the experience is likely forgotten. The process of building memory through repeated recent experiences is "wear-in," while losing it is "wear-out" or "forgetting" (Engel and Blackwell 1982, pp. 249, 483).

We thus operationalize Return Stock as a "stock variable", consistent with previous advertising literature (e.g., Danaher, Bonfrer, and Dhar 2008; Dinner, van Heerde, and Neslin 2014) and the return habit variable used by El Kihal et al. (2025). Return Stock equals previous Return Stock multiplied by a decay parameter to represent forgetting, plus a recent-experience factor incorporating recent returns. The combination of forgetting and recent experience begets wear-out and wear-in respectively. High propensity to forget coupled with only sporadic experiences begets smaller Return Stock.

In sum, Return Stock reflects the connection between past behavior and future actions, making it central to understanding virtuous and vicious cycles. Marketing can set these dynamics in motion starting either with its direct impact on sales or its direct impact on returns. This suggests that marketing can be adjusted – in terms of spending, scheduling and allocation across activities – to leverage returns to drive sales. This increases profits if the gains triggered by the virtuous cycle exceed the costs associated with vicious cycles. A crucial issue for marketing is how strongly it directly influences sales and returns, and how strongly it indirectly influences sales and returns through Return Stock. We will measure this. It requires a dynamic statistical model to capture Return Stock, and a dynamic optimization to recommend marketing spend, timing, and allocation among activities.

Peer Review Version

Research Context and Data

Our research context is online retailing in the fashion industry (clothing, shoes, accessories, etc.). This is a particularly appropriate setting for our study for several reasons.

First, online sales of fashion items represent more than one-third of e-commerce revenues worldwide. It accounted for $668 billion (U.S. dollars) in 2021 worldwide (Statista 2022a), underscoring the importance of this industry to the global economy.

Second, returns data (Statista 2022b) show that clothing and in general fashion items have a significant return rate: 26% in the USA, 32% in the UK, even higher in some European countries (e.g., Austria, Germany, and the Netherlands 33%). The business press documents online retailers trying to find an optimal way to deal with product returns (Ader et al. 2021, Eley 2022). Many retailers (e.g., Asos, Zara, Net-a-Porter) monitor competitors' return policies and consumers' return behavior (Eley 2022).[4] Perhaps ignoring the virtuous cycle, some retailers charge for returns (e.g., Boohoo) or make it onerous for customers to return (e.g., Amazon). Indeed, returns are very relevant to fashion industry management.

Third, fashion retailers spend substantial amounts on price promotions and marketing communications to boost sales. Our framework attests to the importance of marketing in influencing both vicious and virtuous cycles. Marketing will be the decision variable in our optimization.

[4] Our data lack competitive information, limiting our analysis in this respect, but our region-specific findings suggest firms should monitor returns other than their own.

Peer Review Version

Data

We obtained data from a leading retailer that sells fashion items online. Its core products are clothing and shoes for women, men, and kids. It also sells accessories, home décor, and art. It carries a large assortment of brands as well as its private label. It sells worldwide.

Our framework dictates we need data on incidence and monetary size of purchases and returns. When dealing with dynamics, it is useful to analyze consumers' behavior beginning from their first purchase to avoid confounding dynamic and cross section effects. We also need detailed information about the firm's marketing efforts over a relatively long observation period. The retailer provided us with data and information that fulfilled these needs.

We obtained a random sample of new customers acquired in 2013, 2014, or 2015. Each customer is observed from their first purchase till the end of 2017. We also collected demographics (age, gender, country of residence, country of purchase). Additionally, the firm provided its marketing calendar for each country, each day, each year (2013-2017). We organized the information provided in the marketing calendar to create four marketing variables – retailer advertising, brand advertising, markdowns, and free shipping.

We focus on active customers, those who purchased at least once each year after they were acquired. We used this sample for several reasons: (1) *Managerial relevance*: They are a minority of customers but represent 61% of revenues (Table WA.2, Appendix A). (2) *Avoid survivor bias*: Without the "every-year" requirement, we would have included customers who "churned". Customers remaining at the end of our data would be "survivors", different from those we started out with. Thus, what we infer as dynamic behavior could have been cross-customer differences.[5] (3) *Increased estimation precision*. Intuitively, a longer purchase series

[5] We thank the AE for valuable insights on this issue. See Appendix A for discussion suggesting survivor bias might not have been crucial in our application. But at the outset, we thought it prudent to guard against it.

Peer Review Version

provides more *within customer* information for detecting dynamics. Theoretically, work related to initial conditions suggests longer time series improves estimates of dynamics (Keane 1997; Simonov et al 2020). (4) *Consistency with previous research*. Previous work constructs long time series to estimate dynamics (Erdem et al 2003; Mark et al 2013; Seetharaman 2004; Erdem and Keane 1996; Ansari et al 2008; Dubé et al 2010).

The final sample includes 1498 customers and 291,096 observations, which we analyze at the weekly level.[6] Customers live in 78 different countries (52% in Europe, 16% Russia, 15% USA and Canada, 11%, Asia, and 6% in the rest of the world). The return rate, as measured by the percentage of purchases for which at least one item is returned, is 33%. Figure 2 shows return rate differs by country. Figure 3 shows how it trended over time.

Figure 2: Return Rate by Country of Residence.

Note: Y axis is Return *given purchase*, the percentage of purchases where at least one item is returned.

[6] There were 330 customers acquired in 2013 observed for five years (85,800 obs.), 444 acquired in 2014 observed for four years (92,352 obs.), and 724 acquired in 2015 observed for three years (112,944 obs.).

Figure 3: Return Rate over Time.

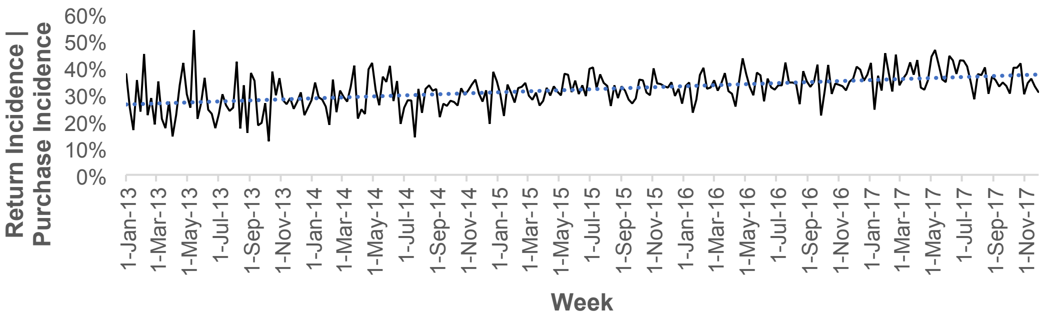

Note: Y-Axis is Return *given* Purchase, i.e., the percentage of purchases where at least one item is returned.

Peer Review Version

Table 2 provides descriptive statistics of the outcome variables of interest: purchase incidence, order size (€) (net of return size), return incidence (whether the customer returned at least one item from the basket purchased in a given week), return size (€) (the value of the items returned) and customer profit per week.[7] So, in 2013, the customers returned at least one item from 26.54% of their purchases, at an average value of €167.21 per return.

Table 2: Descriptive Statistics of Purchase and Return Behavior.

	Purchase Incidence (1=Yes, 0=No)		Order Size per Purchase (net of Return Size)		Return Incidence| Purchase (1=Yes, 0=No)		Return Size per Return		Profits per Customer per Week^A	
Year	**Percentage**	**SD**	**Mean (€)**	**SD**	**Percentage**	**SD**	**Mean (€)**	**SD**	**Mean (€)**	**SD**
2013	8.06	27.22	224.43	216.70	26.54	44.18	167.21	162.60	2.08	45.17
2014	9.00	28.62	205.55	216.99	28.90	45.34	168.37	181.18	5.64	56.81
2015	9.08	28.73	214.12	264.17	32.22	46.74	187.86	243.91	15.18	86.75
2016	10.77	31.00	208.41	265.12	33.70	47.27	186.55	255.92	16.70	92.02
2017	10.25	30.33	209.99	274.79	36.02	48.01	190.50	269.37	15.67	90.51
Tot	*9.77*	*29.70*	*210.69*	*259.90*	33.23	*47.11*	*185.26*	*247.04*	*11.05*	*76.99*

^A Profit calculations follow the methodology described in Equation (8), incorporating actual values of marketing and customer behavior within our sample during the observation period.

[7] In total 33% of customers who purchased in a given year returned at least one item. The average return size is €185, while the average order size is €210. This means customers typically returned part of, not the entire order.

Marketing Activity

The retailer's core marketing activity varies by country and week but not by customer – the retailer did not target marketing at the customer level. The retailer did not advertise on TV, print, or radio. Instead, it reached consumers through display advertising on the retailer's website, its partners' websites, email, and/or social media posts.

We used documents the retailer provided to delineate four types of marketing campaigns: 1) Markdown (MKD), 2) Free shipping (FS), 3) Retailer advertising (RET), and 4) Brand advertising (BRD). Markdown campaigns offered a range of discounts, with reductions between 40% and 90% typically at the end of the season. Free-shipping campaigns didn't follow a consistent timing pattern and were not associated with specific brands but were applicable to the entire order. Retailer advertising aimed to boost the retailer's brand and promote a specific feature (e.g. wish list, private label, sustainable actions room). These campaigns often included shallow discounts (10% or 15% off) co-occurring with holidays or special events (e.g. Mother's Day, preview of the new season). Brand advertising featured manufacturer brands carried by the retailer and did not involve discounts. The retailer funded brand advertising campaigns itself, without co-participation from manufacturers. This reflects the role of the focal firm as an outlet e-retailer that manages its brand assortment and related marketing actions autonomously. Figure 4 provides examples.

Operationalization of the marketing variables

We know the time each campaign begins and ends in each country. Therefore, we see the number of days the retailer ran each campaign type in week t, country c. We operationalize the marketing variables accordingly (see Blattberg and Neslin 1990, p. 368; Mulhern and Leone 1991; p. 70 for a similar approach). For example, let week t be the week of Monday, February 6 thru Sunday, February 12, and c be the USA. If the retailer ran a retailer advertising campaign

Peer Review Version

Figure 4: Examples of Marketing Activities.

Panel A: Brand-Building

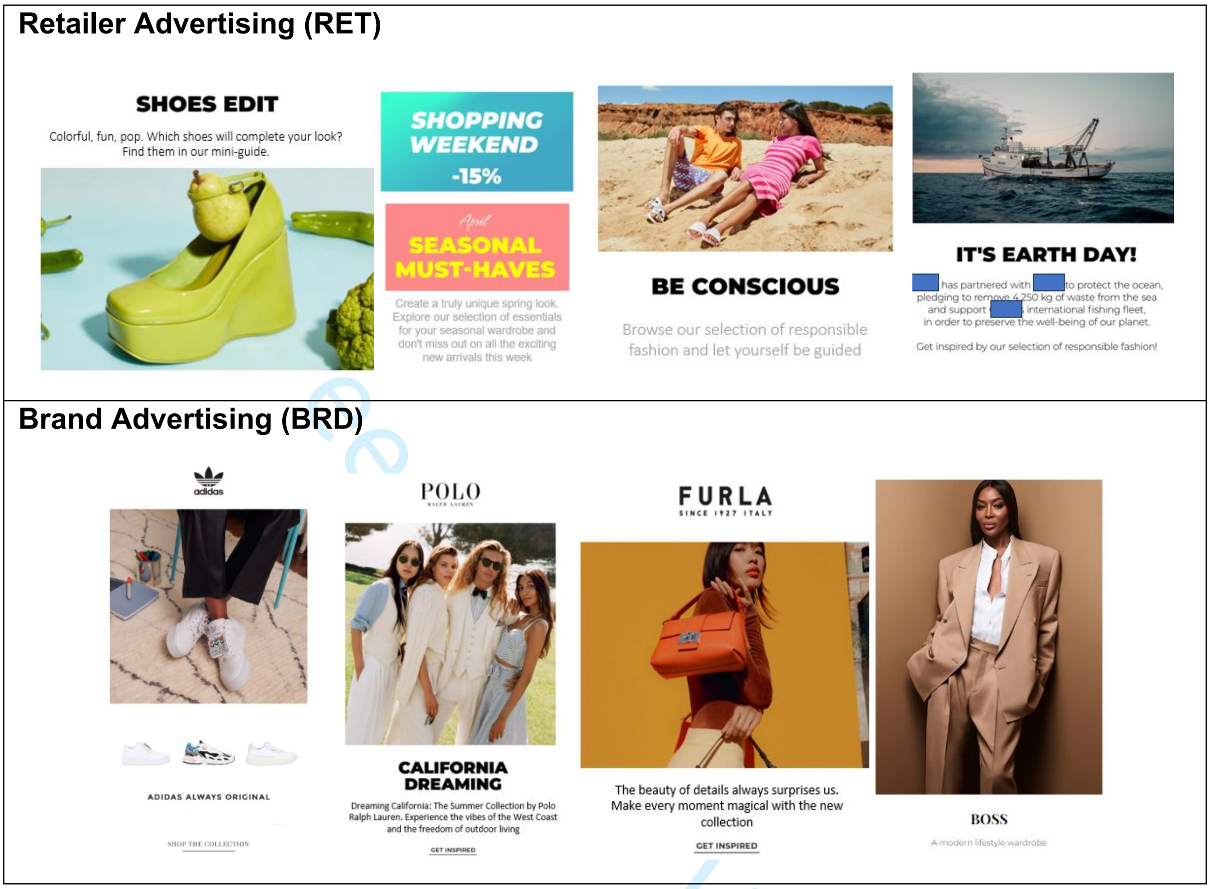

Panel B: Non-Brand Building

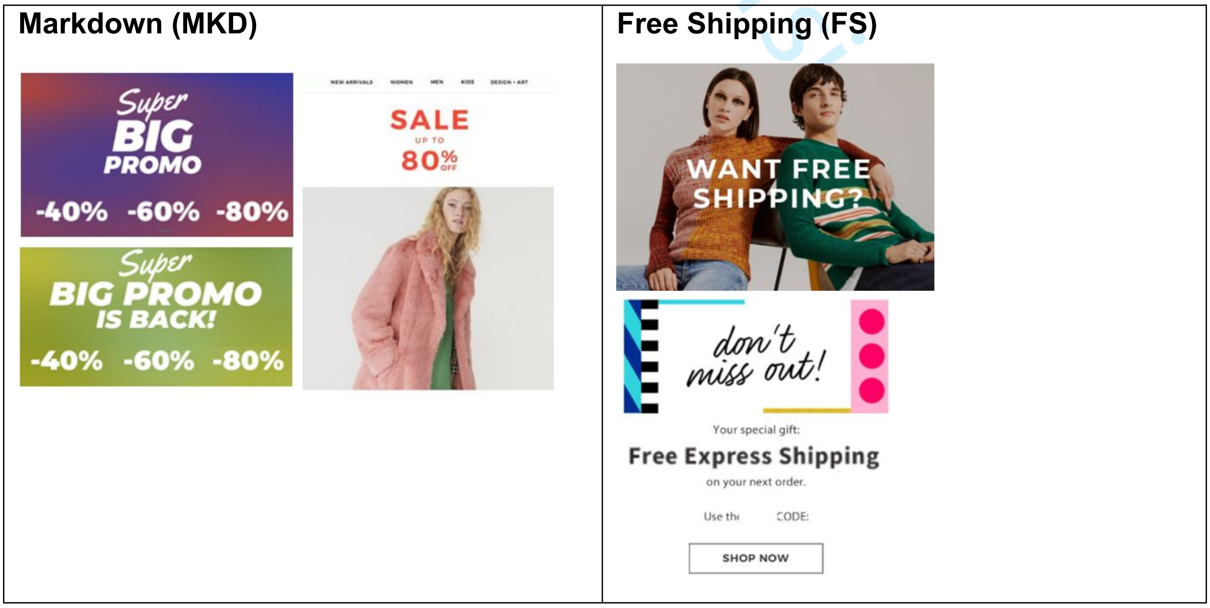

Note: We removed the name of the retailer from some examples.

in the USA between February 6 and February 8 and ran a free shipping campaign from February 7 till Friday, February 10, retailer advertising (RET) would equal three days in week *t*, country *c* (USA) and free shipping (FS) would equal four days in week *t*, country *c*. Table 3 summarizes the marketing variables.

Table 3: Frequency of Marketing Activity.

Panel A: Non-Brand Building.

	Discounts (MKD)[d]						Free Shipping (FS)					
Year	**M[a]**	**SD**	**% Non-Zero[b,c]**	**Cond. M[d]**	**Min**	**Max**	**M[a]**	**SD**	**% Non-Zero[b]**	**Cond. M[d]**	**Min**	**Max**
2013	0.47	0.62	40.47	1.17	0	2	0.07	0.43	3.00	2.34	0	3
2014	0.92	1.82	41.88	2.20	0	7	0.08	0.45	3.28	2.44	0	3
2015	0.35	0.60	30.88	1.14	0	5	0.09	0.46	3.29	2.73	0	3
2016	0.33	0.57	28.83	1.13	0	4	0.08	0.45	3.24	2.47	0	3
2017	0.30	0.47	30.08	1.01	0	7	0.10	0.47	3.54	2.82	0	3
Total	0.48	0.82	34.43	1.33	0	7	0.09	0.46	3.21	2.80	0	3

Panel B: Brand Building.

	Manufacturer Brand (BRD)						Retailer Brand (RET)					
Year	**M[a]**	**SD**	**% Non-Zero[b]**	**Cond. M[d]**	**Min**	**Max**	**M[a]**	**SD**	**% Non-Zero[b]**	**Cond. M[d]**	**Min**	**Max**
2013	0.53	0.60	47.61	1.12	0	2	1.37	2.28	45.57	3.00	0	7
2014	3.43	2.97	67.53	5.08	0	8	0.65	1.84	12.66	5.13	0	8
2015	1.87	2.34	50.33	3.71	0	7	0.64	1.83	12.43	5.13	0	8
2016	3.37	1.82	90.69	3.72	0	7	0.64	1.83	12.41	5.16	0	8
2017	3.38	1.70	92.12	3.67	0	7	0.65	1.85	12.58	5.17	0	8
Total	2.52	1.88	69.66	3.46	0	8	0.79	1.93	19.13	4.72	0	8

[a] M indicates the mean number of days per week each marketing activity was implemented, across countries.

[b] % Non-Zero indicates the percentage of weeks with non-zero marketing.

[c] The percentage discount for MKD varies with deep discounts over 70%, usually at the end of the season, occurring in 8.17% of cases.

[d] Cond. M indicates the mean conditional on the presence of the marketing campaign.

Since marketing is operationalized as the number of days the retailer ran each campaign type in week *t* and in country *c*, the minimum of each marketing variable was 0. The maximum could be greater than 7. It is possible the retailer ran two or more campaigns of the same type in the same week. For example, the retailer might run two brand-building campaigns, each for a different brand and each for five days, in week *t*. BRD_t would equal 10 days. For example, in 2017 the retailer offered markdowns (MKD) on average .30 days per week, with a maximum of 7, and retailer brand-building campaigns (RET) an average of .65 times per week, with a maximum of 8.

Statistical Model

The DSS conceptual model, our Framework in Figure 1, suggests four customer decisions: 1) whether to purchase in a given week (purchase incidence), 2) how much to spend (order size) if purchase, 3) whether to return at least one item from the order purchased in a given week (return incidence), 4) how much to return (return size) if return. Accordingly, we have a statistical model consisting of four equations. Purchase incidence is conditional on positive latent utility for purchase. Order size is observed only if purchase. Return incidence is conditional on purchase and on positive latent utility for return. Return size is observed only if purchase and return. Purchase and return incidence are probit models. Order and return size are regressions using the logarithm of the dependent variables. The four equations are:

<u>Purchase</u>

$$P_{it} = \begin{cases} Purchase & if \quad P_{it}^{*} > 0 \\ No\ Purchase & if \quad otherwise \end{cases} \tag{1.a}$$

where

$$P_{it}^{*} = \alpha_i^P + \gamma^P ReturnStock_{it} + \beta_M^P Mktg_{ct} + \beta_{MLag}^P Mktg_{ct-1} +$$

$$\iota^P P_{it-1} + \kappa^P CC_i + \nu^P T_t + u_{it}^P \tag{1.b}$$

Peer Review Version

Order Size

$$OS_{it} = \alpha_i^{OS} + \gamma^{OS} ReturnStock_{it} + \beta_M^{OS} Mktg_{ct} + \beta_{MLag}^{OS} Mktg_{ct-1} + \iota^{OS} OS_{it-1} +$$

$$\kappa^{OS} CC_i + \nu^{OS} T_t + u_{it}^{OS} \quad if\ Purchase;\ else\ unobserved \tag{2}$$

Return Incidence

$$R_{it} = \begin{cases} Return & if \quad R_{it}^* > 0 \\ No\ Return & if \quad otherwise \end{cases} \tag{3.a}$$

where

$$R_{it}^* = \alpha_i^R + \gamma^R ReturnStock_{it} + \beta_M^R Mktg_{ct} + \beta_{MLag}^R Mktg_{ct-1} + \tau^R OS_{it} + \kappa^R CC_i +$$

$$\nu^R T_t + u_{it}^R \quad if\ Purchase;else\ unobserved \tag{3.b}$$

Return Size

$$RS_{it} = \alpha_i^{RS} + \gamma^{RS} ReturnStock_{it} + \beta_M^{RS} Mktg_{ct} + \beta_{MLag}^{RS} Mktg_{ct-1} + \iota^{RS} RS_{it-1} + \tau^{RS} OS_{it} +$$

$$\kappa^{RS} CC_i + \nu^{RS} T_t + u_{it}^{RS} \quad if\ Purchase\ and\ Return;else\ unobserved \tag{4}$$

As suggested by our DSS-motivated conceptual model (Figure 1), $ReturnStock_{it}$ represents consumer *i*'s memory regarding returns. It draws on "AdStock" (Broadbent 1984, Danaher, Bonfrer, and Dhar 2008). Define $Return\ Share_{it-1}$ as the ratio of return size (RS_{it-1}) to order size (OS_{it-1}) – the fraction of an order that gets returned. λ captures wear-in and forgetting parsimoniously. Large λ increases wear-in and inhibits forgetting. A vivid experience will boost memory significantly and not be forgotten easily. Using return share is consistent with the theory of interference in memory research. It suggests that returning only part of an order leaves a weaker memory trace than full returns, due to competing information that interferes with memory encoding and retrieval (Kliegl and Bäuml 2021). Return Stock is then:

$$ReturnStock_{it} = \lambda ReturnStock_{it-1} + Return\ Share_{it-1} \tag{5}$$

Peer Review Version

The parameters γ^P and γ^{OS} in Equations 1 and 2 quantify the virtuous cycle, capturing the impact of Return Stock on purchase. The parameters γ^R and γ^{RS} in Equations 3 and 4 capture the vicious cycle, quantifying the impact of Return Stock on returns. Per our framework, we include the direct impact of marketing on both purchase and returns by including $Mktg_{ct}$ in all four equations. $Mktg_{ct}$ is a 4-element vector of the four marketing types– retailer brand advertising, brand advertising, free shipping and markdown. The subscripts *c* and *t* indicate that marketing activity varies by country and week.

We include several control variables. Lagged values of the dependent variable in Equations 1, 2 and 4 (Return Stock includes lagged returns, so we do not need to add a lagged term in Equation 3) reflect carryover in the dependent variable separate from other dynamics, notably $ReturnStock_{it}$. CC_i are customer *i*'s descriptors. They include gender, age, country of residence, and acquisition cohort, as well was past behavior such as the number of purchases, average order size, return share, and average return size given return, and the proportion of purchases in each channel (tablet, smartphone, or desktop) in the 26 weeks prior to the beginning of the observation window for customer *i*.[8] See details in Web Appendix A. The four equations also control for time trend (T_t), using fixed effects for month and year of purchase. Finally, we include current order size OS_{it} in Equations 3 and 4, the assumption being that larger orders are more likely to have at least some items returned, and naturally the monetary size of those returns will be larger if larger orders are returned.

The error terms follow a multivariate normal distribution $u_{it}^P, u_{it}^{OS}, u_{it}^R, u_{it}^{RS} \sim N(0, \Sigma)$. The covariance matrix Σ among the error terms is:

[8] This is for initialization. The first 26 weeks are not included in the estimation, yielding 252,148 observations instead of the 291,096 available.

Peer Review Version

$$\Sigma = \begin{bmatrix} 1 & \sigma_{OS,P} & \sigma_{R,P} & \sigma_{P,RS} \\ \sigma_{P,OS} & \sigma_{OS}^2 & \sigma_{R,OS} & \sigma_{OS,RS} \\ \sigma_{P,R} & \sigma_{OS,R} & 1 & \sigma_{R,RS} \\ \sigma_{P,RS} & \sigma_{OS,RS} & \sigma_{R,RS} & \sigma_{RS}^2 \end{bmatrix} \quad (6)$$

As is common, we set the variance of the purchase and return probit model equations to 1.

We identify the parameters for OS_{it} in Equations 3 – 4 using exclusion restrictions. The CC_i variables differ across equations: Equation 1 includes total purchase incidence; Equation 2 includes average order size. Equations 3 and 4 do not. Lagged order size OS_{it-1} should influence current order size OS_{it} but not directly affect current return incidence R_{it} or return size RS_{it}. Therefore, it is in Equations 1 and 2 but not 3 or 4. The lagged dependent variable also differs across equations.

We leverage customer-level panel data to identify the impact of Return Stock. Recall our sample is customers who bought each year starting the year they were acquired. This contributes to identification as follows: First, *temporal variation*. $ReturnStock_{it}$ captures customer *i*'s return experience up to time *t*. Since subsequent behavior is observed after *t*, this reveals the impact of prior returns on future behavior, on a within-customer basis.[9] We also include customer-specific temporal controls such as lagged return and purchase. Second, *theoretical evidence.* Consistent with intuition, econometric work on initial conditions for data required by dynamic models suggests that more complete time series improve identification of dynamic coefficients (Keane 2003; Simonov et al 2020). Simonov et al state, "…only when we use larger values of T^{purch} [do] we see the true values [of a loyalty effect]" (p. 797).[10]

[9] This also enables us to validate our results using propensity score matching (PSM, see Web Appendix F for more details). The close alignment between PSM and model-based results provides further supports our results.

[10] The authors gratefully acknowledge Asim Ansari (Columbia University) for his insights and suggestions on this issue.

Peer Review Version

These factors require sufficient variation within customer over time in return incidence and size. The data provide this. For example, only 26% of customers never returned a purchase; 2% returned all their purchases. This means 72% returned at least once but not always. Among these, the mean return rate was 37% per customer - maximum was 93%; minimum was 4%. In terms of the return size, among customers who returned at least twice, 0% returned the same amount for each return. So, return size varies within customer. The mean return size per customer was €153 - max was €1407; min was €17. These return rate and amount numbers suggest adequate variation *within customers* to measure the impact of returns.

A third factor aiding identification is *cross-country variation*. Countries vary in return policies – strict versus lenient. In addition, countries differ in the timing of marketing, creating temporal variation. Thus, variation across countries systematically generates temporal and cross-sectional variation.

The above three factors support our claim of parameter identification. We have much variation in the data for identifying dynamics. However, we cannot guarantee causal inference. For example, we cannot capture how consumers' valuations of returns evolve over time, nor do we observe their opinions of the items they receive. These unobservables may influence subsequent purchase and return behavior. This limitation is consistent with the DSS viewpoint, which emphasizes insights as well as concrete policy. The policies may not be exact because parameters are not quantified perfectly; however the insights, while judgmental, remain.

We estimate the model jointly using the simulated maximum likelihood SUR procedure in Stata's CMP package (Roodman 2011, p. 160). This is full-information maximum likelihood (FIML). It considers the joint likelihood function of all equations, enabling consistent estimation. (Roodman 2011, Lahiri and Schmidt 1978).

Peer Review Version

Model Optimization

A key step in the DSS paradigm is optimization – in our case, prescribing the marketing policy that exploits the virtuous and vicious cycles produced by product returns. The estimated statistical model provides the results required for this step. We need to derive the 4-element vector $Mktg_t$,[11] the number of days in week *t* for each marketing type (discounts (*MKD*), free shipping (*FS*), retailer advertising (*RET*), and brand advertising (*BRD*)). Each campaign has its own parameter in Equations 1 – 4. We optimize *N* customers over a finite horizon W=25 weeks. We initialized $ReturnStock_{it}$ at 0.75.[12] Insights were not sensitive to the initial value. The optimization is:

$$\max_{Mktg_t} \sum_{i=1}^{N} \sum_{t=1}^{W} Profit_{it} \tag{7}$$

$$Profit_{it} = P_{it} \times Margin \times \{OS_{it} - R_{it} \times [RS_{it} + Processing_{Costs t}]\} -$$

$$MKD\_Spend_t + FS\_Spend_t + RET\_Spend_t + BRD\_Spend_t \tag{8}$$

$Profit_{it}$ is profit for customer *i*; P_{it} is the probability customer *i* purchases in week *t*. We use an average *Margin* provided by the retailer (40%) that multiplies times revenue. Revenue depends on order size and returns costs. Our order size data (*OS*) include any price discount or free shipping cost. Returns costs equal the probability of a return (*R*) times the return amount (*RS*) plus processing costs. Return size is customer-specific, predicted by our model; however the processing cost is the same for each customer, so it is just subscripted by *t*. Management told us returns costs were the same each week, 8€, reflecting processing and return shipping.

---

[11] To simplify, in this optimization we do not distinguish marketing by country, so we drop the *c* subscript. We optimize for specific countries in Web Appendix E.

[12] We set the initial value of *ReturnStock* via an iterative search procedure. Using a 25-week horizon, we explored initial levels from 0.45 to 0.75 (step 0.10) and chose the smallest value that yielded stable process-behavior (control-chart) trajectories for *ReturnStock*; 0.75 was the first to satisfy this criterion.

Peer Review Version

The weekly spend for each campaign type equals the number of days that campaign type is run in week t times its cost per day. The retailer used the same 70-cent daily cost for each campaign type. This cost includes two main elements: i) copy and creative design, ii) opportunity cost of house ads appearing on the retailer's website – website space is precious; using it for one activity precludes using it for another. Marketing spend is not contingent on purchase because as noted above, purchase-related costs such as discounts and the cost of free shipping are included in order size OS_{it}.

Overall, the cost structure means our decision variables are the number of days to run each campaign type in week t, with a cost of 70 cents per day for each campaign type – the sum of costs across the four campaign equals spend. [13]

The optimization contains immediate and long-term factors. First, Figure 1 shows that marketing increases purchases and returns directly. It thus immediately provides a benefit (more sales) and either a benefit (if it directly decreases returns) or a cost (if it directly increases returns). Second are the dynamics. Returns in period t directly determine purchase as well as returns in future periods (through either the virtuous or vicious cycles), and purchases in period t also determine purchases and returns in future periods.

How this plays out in optimal levels of marketing depends on the estimated statistical model parameters. But one can see the necessity of dynamic optimization to resolve short- and long-term tradeoffs. For example, marketing may have a hugely positive impact in the current period if it increases purchases and does not have a positive direct effect on returns. But more purchases will beget more returns through Equation 3, and those returns may increase purchase

[13] When optimizing, we set constraints for marketing decision variables based on the range of data in Table 2 and practical considerations.

Peer Review Version

in the future due to the virtuous cycle However, we also have a vicious cycle – current returns increase Return Stock and increase returns in the future.

The optimization must determine weekly marketing spend that accounts for both virtuous and vicious cycles. Our model quantifies these relationships, but this is a dynamic optimization, not the sort of calculation one makes "back of the envelope". We will test whether dynamic optimization improves over myopic optimization.

Customer-level purchases and returns are probabilistic, so we require a stochastic dynamic optimization. Since the firm used a one-quarter time horizon when specifying its marketing, we employ a finite horizon (Khan, Lewis, and Singh 2009). We use the GRG algorithm implemented in Frontline Solver to optimize. The GRG procedure is a nonlinear optimization algorithm. When it converges to a solution, it means that there is no other set of values for the decision variable *close* to the ones identified that yield higher profits over our time horizon; however, there may be better solutions far away from the one identified. We address the possibility that the algorithm identified a local optimum using multiple starting values (See Web Appendix C for more details on optimization specification).

As noted earlier, we use myopic optimization as a benchmark to compare dynamic versus myopic (see Neslin et al. 2009, pp.728-729). Our myopic optimization uses a simple rule: find the optimal level of marketing activity each week without considering its implications for future marketing activities, which means considering only the current marketing efforts at time *t*. As Neslin et al. report and is sometimes the case, myopic optimization, a form of "greedy algorithm", can do quite well, especially if the effect sizes of dynamics are not very large.

Peer Review Version

Results

Estimates of the Purchase / Returns Model

Return Stock (Equation 5) captures the magnitude of each customer's memory of past returns. Its coefficients in the statistical model represent its association with purchases and returns, i.e., the vicious and virtuous cycles. We performed a grid search to estimate λ, letting the value range from 0 to 1 with increments of .1 (Dinner, van Heerde, and Neslin 2014). Our first pass suggested we explore the area above .7 more granularly. Figure 5 shows the results, suggesting a peak between .9 and .95. We then narrowed the range and conducted another grid search using increments of .01. Figure 5 shows a clear peak at $\lambda = .93$. This suggests using .93 as the estimate of λ (see Fader, Lattin, and Little 1992). This large value suggests consumers remember return experiences a long time.[14] The values of the coefficients for Return Stock measure the impact of this memory on returns and purchase.

Figure 5: Estimating $ReturnStock_{it}$ Carryover: Model Fit as a Function of λ.

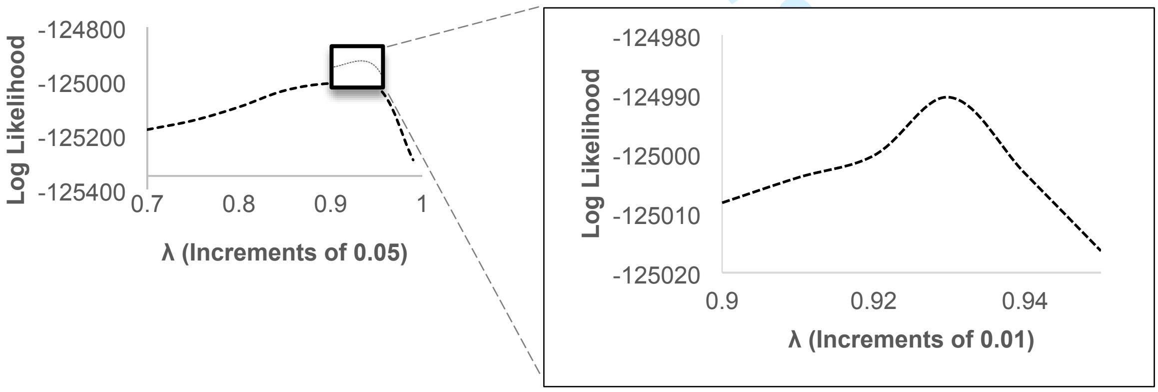

[14] Calculating Equation 5 over time shows that with λ = .93, 90% of the impact of returns occurs within 8½ months, which seems reasonable given the relatively infrequent interpurchase time for fashion clothing.

Peer Review Version

Table 4 presents the estimates of the proposed statistical model, Equations 1 – 4 (Model 1). These confirm that returns are associated with future sales, as $ReturnStock_{it}$ has a significant and positive impact on both purchase incidence (γ^{P} = .226, p=.000) and order size (γ^{OS} = .118, p=.000). This captures the virtuous cycle. Table 4 also reveals a vicious cycle – $ReturnStock_{it}$ has a positive impact on return incidence (γ^{R} = .249, p=.000) and return size (γ^{RS} = .090 p=.000). Returns increase future sales but generate more future returns as well.

Both non-brand-building (β^{P}_{MKD} = .010, p=.000, β^{P}_{FS} = .143, p=.025) and brand-building marketing (β^{P}_{RET} = .014, p=.000, β^{P}_{BRD} = .102, p=.000) are positively and significantly associated with purchase. Free shipping and retailer advertising also are associated with order size (β^{os}_{FS} = .110, p=.000, β^{os}_{RET}= .006, p=.003). Lagged marketing discounts are associated with purchase incidence and order size ($\beta^{P}_{MKD.LAG} = .016, \text{p} = .000, \beta^{os}_{MKD.LAG}$ = .017, p=.002). This may reflect consumers waiting for deeper discounts.[15] Negative lags could reflect purchase acceleration – purchases move forward in time. Free shipping promotion is associated with purchase and order size and is the only marketing associated with return incidence (β^{R}_{FS} = .092, p=.053), supporting the power of these promotions (Shehu, Papies, and Neslin 2020).

We further test for dynamic effects of product returns by comparing the proposed Model 1 (M1; Equations 1– 4) with two benchmarks: Model 2 (M2) limits returns dynamics by imposing λ = 0 in computing $ReturnStock_{it}$; Model 3 (M3) eliminates returns dynamics by excluding $ReturnStock_{it}$ and RS_{it-1} from Equations 1 – 4. Results are in Table W.B2. They confirm the proposed model (M1) is superior to both M2 and M3 in terms of AIC and BIC. This supports a dynamic relationship between product returns and sales.

---

[15] In fact, the firm typically increases discount rates progressively (e.g., from 60% in the first week to 80% in the second week). Consumers adjust their behavior for this (Gonul and Srinivasan 1996).

Peer Review Version

Table 4: Model 1 (M1) - Proposed Model – Full Product Returns Dynamics.

Macro Category	**Variable**	**Purchase Incidence (P_{it})** Coef. (SE)	**Order Size (OS_{it})** Coef. (SE)	**Return Incidence (R_{it})** Coef. (SE)	**Return Size (RS_{it})** Coef. (SE)
Marketing: Non-Brand Building	Marketing Discounts (β_{MKD})	0.010	0.006	-0.024	-0.012
		(0.004)	(0.005)	(0.012)	(0.008)
	Free Shipping (β_{FS})	0.143	0.110	0.092	0.028
		(0.019)	(0.023)	(0.048)	(0.028)
Marketing: Brand-Building	Retailer Advertising (β_{RET})	0.014	0.006	-0.002	0.000
		(0.002)	(0.002)	(0.005)	(0.003)
	Brand Advertising (β_{BRD})	0.102	-0.030	0.013	0.007
		(0.028)	(0.032)	(0.066)	(0.045)
Lagged Marketing: Non-Brand Building	Marketing Discounts Lag ($\beta_{MKD.LAG}$)	0.016	0.017	0.012	-0.001
		(0.004)	(0.005)	(0.012)	(0.008)
	Free Shipping Lag ($\beta_{FS.LAG}$)	-0.022	-0.021	0.001	0.032
		(0.020)	(0.025)	(0.053)	(0.028)
Lagged Marketing: Brand-Building	Retailer Advertising Lag ($\beta_{RET.LAG}$)	-0.005	-0.004	0.005	0.001
		(0.002)	(0.002)	(0.005)	(0.003)
	Brand Advertising Lag ($\beta_{BRD.LAG}$)	-0.064	-0.053	0.031	-0.098
		(0.029)	(0.033)	(0.077)	(0.067)
Product Returns Dynamics	Return Stock (γ)	0.226	0.118	0.249	0.090
		(0.007)	(0.015)	(0.020)	(0.009)
State Dependence	Purchase Incidence Lag (ι^{P})	0.320			
		(0.014)			
	Order Size Lag(ι^{OS})		0.082		
			(0.004)		
	Return Size lag (ι^{RS})				0.737
					(0.025)
Recursive Component	Order Size (τ)			0.325	0.042
				(0.061)	(0.004)

Notes: Number of observations= 252148, LL = -124990.48, Wald χ^2 = 16513.50, $p < .001$
Order Size and *Return Size* variables are expressed as natural logarithm
κ^P, κ^R, κ^{OS}, κ^{RS} parameters for control variables and ν^P, ν^R, ν^{OS}, ν^{RS} for time trend variables are not reported.
Complete results are available in the Web Appendix B (Table W.B1)

In summary, the statistical model reveals 1) a virtuous cycle reflected in a positive association between returns and future purchases, 2) a vicious cycle reflected in a positive association between returns and future returns, 3) a positive relationship between marketing and purchase, 4) a positive relationship between one marketing activity, free shipping, and return incidence, and 5) a model including return dynamics improves fit over models that do not.

Peer Review Version

Determining how to free-ride the positive association of returns with purchase is a complex task given its dynamic associations with purchase incidence, order size, return incidence, and return size. It requires marketing to capitalize on the virtuous cycle. This can entail adjustment of: (1) marketing spend, (2) schedule of marketing activities, and (3) spend allocation across marketing efforts. We present results along these three dimensions.

Free-Riding Product Returns: Optimal Marketing Spend

Table 5 reports dynamic optimization increases profits by 46.56% in comparison to myopic optimization.[16] It decreases both marketing spend (39.87% lower) and returns costs (33.18% fewer returns at close to the same monetary value per return).

Table 5: The Full Dynamic Model (M1) Optimized Either Dynamically or Myopically: Average Sales, Returns, Profits and Marketing Costs per Customer per Week.

	(M1 / Myopic) (€)	**(M1 / Dynamic) (€)**	**% Difference**
Sales	126.81	139.77	10.22
Marketing Spend	15.08	9.07	-39.87
Returns	0.40	0.27	-33.18
Returns Monetary Value	15.66	15.85	1.23
Profits	26.17	38.36	46.56

Note: Figures represent averages calculated across all customers and over weeks 4 – 25.

Table 5 optimizes using the dynamic statistical model M1. However, maybe this isn't necessary. Perhaps limiting or eliminating returns dynamics – M2 and M3 – would generate similar profits. M1's complex modeling of dynamics may not be required.

To investigate, we compute the optimal marketing policy dictated by M1, M2, and M3, using dynamic as well as myopic optimization. This yields six marketing policies: *(M1/ Myopic)*, *(M1/Dynamic)*, *(M2/Myopic)*, etc. M1 is the "correct" model as it fits the data best. We therefore substitute the marketing plan prescribed by each of the six policies into M1 to

[16] The first periods of a dynamic optimization are often transitory. Stability for our optimization began in the fifth period ("t_4"). We therefore state profits and comparisons from t_4 through t_{25}. We call this the "evaluation period".

Peer Review Version

compute profits. The objective is to show how much one would lose by using the "wrong" statistical model / optimization combination.

Figure 6 shows profits for the six policies. The two bars on the far right are from Table 5 – dynamic optimization of M1 generates more profit than myopic optimization. Figure 6 adds two insights: First, dynamic optimization is better than myopic *as long as the statistical model includes dynamics*: (*M1/Dynamic)* beats *(M1/Myopic)*, *(M2/Dynamic)* beats *(M2/Myopic)*, but (*M3/Dynamic)* does not beat (*M3/Myopic). M3* omits *all* returns dynamics, so dynamic optimization has no leverage. Second, dynamic optimization generates more profit the closer the statistical model gets to the fully dynamic *M1* (the black bars in Figure 6); myopic optimization does not (the gray bars). Myopic optimization does not consider dynamics, so does not exploit a statistical model that includes dynamics. Overall, Figure 6 shows it is important to include dynamics *both* in the statistical model *and* in the optimization.

Figure 6: Optimal Profits Generated by Each Descriptive Model / Optimization Combination.

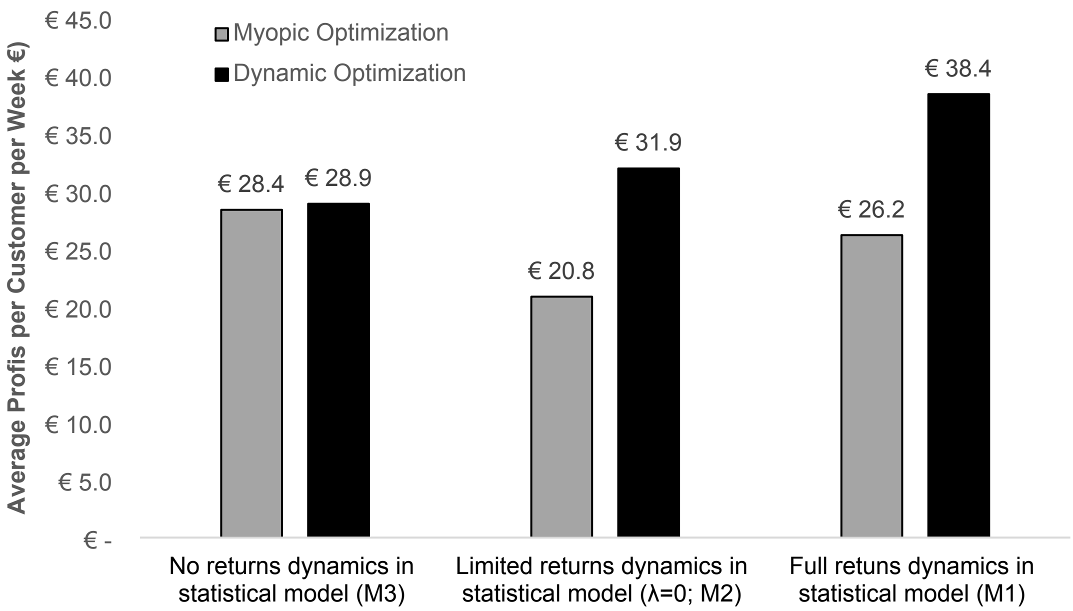

Notes:
M1: Full Model (*$ReturnStock_{it}$*): The full proposed model described in Equations 1-6
M2: *Use λ=0 to compute $ReturnStock_{it}$*: Limits return dynamics by using λ=0 in calculating *$ReturnStock_{it}$*
M3: *No Return Dynamics*: Completely eliminates return dynamics by excluding *$ReturnStock_{it}$* and *RS_{it-1}* from Equations 1-6

Free-Riding Product Returns: Optimal Marketing Scheduling

Figure 7 shows the marketing spend and returns patterns for (*M1/Dynamic*).[17] It shows a pulsing pattern for marketing and a pulsing pattern of Return Stock one week out of phase with the spend pattern. A marketing pulse increases purchases. Some are returned, increasing Return Stock in the next week. We can free-ride this increase because larger Return Stock is virtuous. It stimulates purchase, so we can cut back marketing when Return Stock increases. The crucial insight is firms can increase profits by using marketing to exploit the positive effects of returns – the virtuous cycle.

Figure 7: Dynamic Optimization: Average $ReturnStock_{it}$ and Marketing Spend

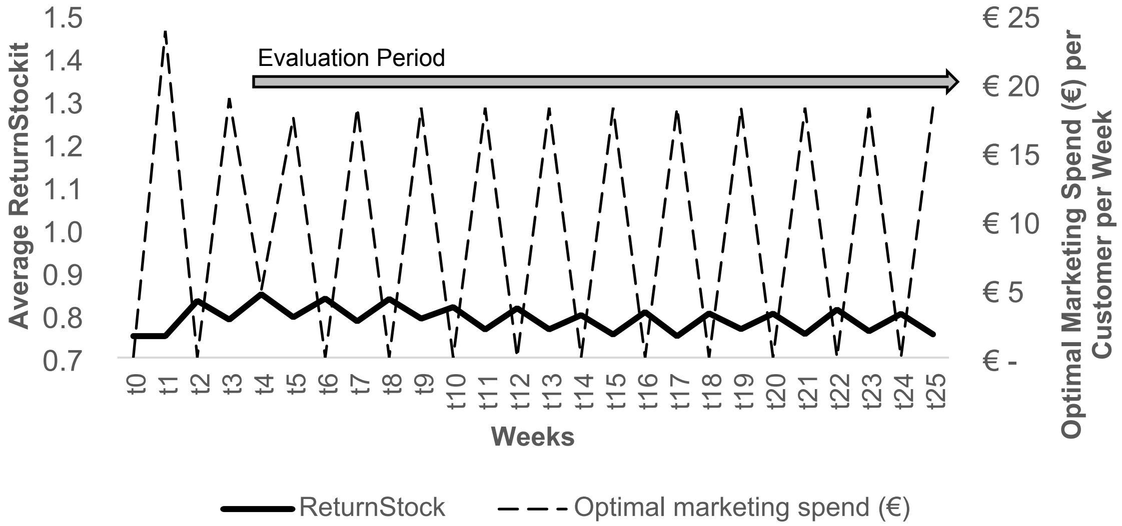

Note: The first periods of dynamic optimization are often transitory in approaching a stable solution. The progression of the key variable Return Stock suggests stability for our optimization began in the fifth period ("t_4"). We therefore state all profits and profit comparisons from t_4 through t_{25} - the "evaluation period" in Figures 7-8.

Figure 8 myopically optimizes *M1*. It shows a hint of pulsing but not nearly what we see with dynamic optimization. Myopic optimization does not *capitalize on* the virtuous cycle because it does not consider future benefits of current actions.

[17] Profit calculations are from t_4 through t_{25}. Also, note t_4 begins with marketing spend not hitting 0 (Figure 7) and myopic showing a "bump" (Figure 8). This is because both optimizations used MKD in t_4.

Figure 8: Myopic Optimization: Average $ReturnStock_{it}$ and Marketing Spend.

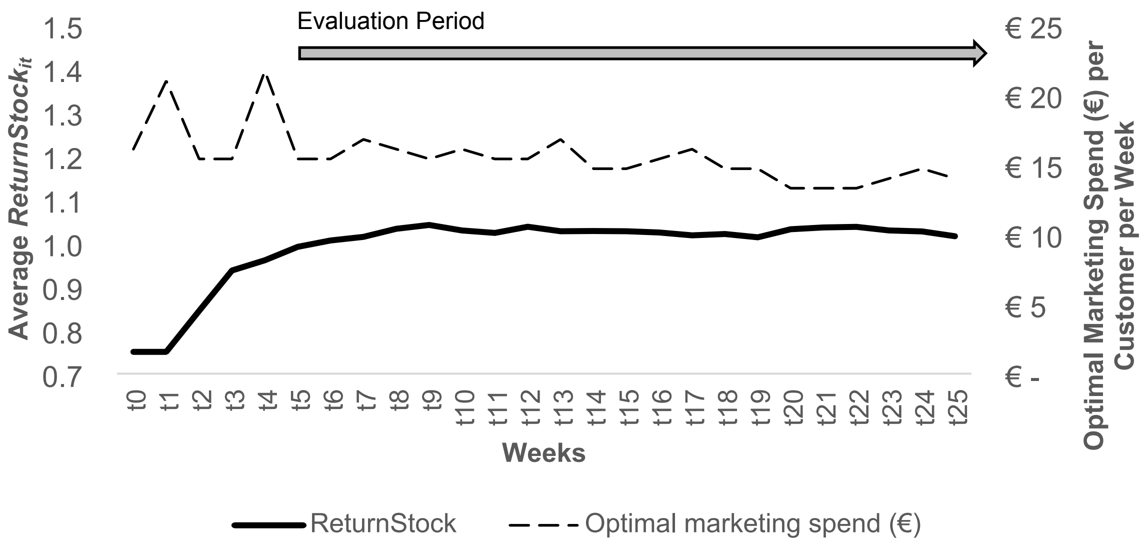

Free-Riding Product Returns: Optimal Marketing Allocation

The success of free-riding in Figure 6 entails scheduling – a pulsing strategy for *(M1/Dynamic)*. Table 6 shows the role played by reallocating among marketing activities. Most notable is that dynamic optimization allocates 22.95% of its marketing to free shipping and 37.67% to retailer advertising. In contrast, myopic optimization allocates 28.22% to free shipping; 32.99% to retailer advertising.

Table 6: Marketing Spend Allocation: Full Dynamic Model (M1) Optimized Either Dynamically or Myopically.

		Optimal Average Marketing Spend (€) per Customer per Week		**Allocation Proportion (%)**		**Comparison in Allocation Proportion (%)**
Type of Marketing Activity		Myopic Opt.	Dynamic Opt.	Myopic Opt.	Dynamic Opt.	Dynamic vs. Myopic
Non Brand-Building	Marketing Discounts	0.22	0.22	1.45	2.40	0.94[a]
	Free Shipping	4.33	2.13	28.22	22.95	-5.27
Brand - Building	Retailer Advertising	5.06	3.50	32.99	37.67	4.68
	Brand Advertising	5.73	3.44	37.34	36.99	-0.36
Total		15.34	9.29	100	100.00	100.00

a. 0.94% ≈ 100% × (2.40/1.45 – 1) (2.40/1.45 – 1 exactly equals 0.66; difference is due to rounding.

Review Version

The parameter estimates in Table 4 suggest why this happens. It shows free shipping is strongly associated with purchase incidence and order size, initiating the virtuous cycle. However, it also stimulates returns. Some of these returns contribute to the virtuous cycle, but they also stimulate future returns, the vicious cycle. The myopic optimization doesn't see this, so happily allocates more of its budget to free shipping. Dynamic optimization in contrast recognizes the vicious cycle, so it pulls back from free shipping.

Dynamic optimization also realizes that retailer advertising relates positively to both purchase incidence and order size, the only marketing other than free shipping that does this. However, retailer advertising doesn't directly increase returns, hence one can allocate it more funds to generate the virtuous cycle without worrying about generating too many returns. It takes forward looking optimization to recognize this, which myopic doesn't do.

Again we see dynamic optimization works because it incorporates the future into its calculations. It therefore allocates funds among marketing activities to leverage the virtuous cycle without creating too much of the vicious cycle.

Quantifying Free-Riding Opportunities Across Geographical Regions and Return Policies

The focal firm operated in several countries. Due to country-specific competition and culture, the firm's return policy differed by country. The USA and Asia/Australia offered lenient policies entailing free returns and up to 100 days between purchase and a permissible return. Russia and European countries such as Denmark, Greece, and the Netherlands had strict return policies with payment required for returns and short return windows (e.g., 15 days). This provides an opportunity to quantify the free-riding opportunity by examining the impact of lenient *versus* strict return policies.

Our analysis consists of the following steps for each country: (1) Estimate the fully dynamic statistical model, *M1*. (2) Calculate the virtuous and vicious cycles using *M1*. (3)

Peer Review Version

Calculate "delta", the difference between virtuous and vicious cycles. (4) Optimize marketing using *M1* with both dynamic and myopic optimization. The difference between dynamic and myopic optimization indicates free-riding. (5) Relate this difference to delta.

Bower and Maxham (2012) and Wood (2001) suggest lenient return policies increase the virtuous cycle so delta will be highly positive for countries with lenient policies. This suggests lenient policies are associated with more free-riding through the following logical chain: large delta (virtuous minus vicious cycles) => large difference in dynamic minus myopic optimization => free-riding. We investigate by going through the steps outlined above.

Step 1: We estimate *M1* for each of four geographical "regions": two representing lenient policies (USA, Asia & Australia), and two representing strict policies (European Countries listed in Table W.E1, and Russia). Results are detailed in Web Appendix E.[18]

Step 2: We calculate revenues (incidence × order size) when $ReturnStock_{it}$ is at its mean plus one standard deviation, minus revenues when $ReturnStock_{it}$ is at its mean. This is the virtuous cycle. The vicious cycle is returns costs (return incidence × return size) when $ReturnStock_{it}$ is at its mean plus one standard deviation, minus costs when $ReturnStock_{it}$ is at its mean. Web Appendix D shows details.

Step 3: The left y-axis in Figure 9 graphs delta, the difference between virtuous and vicious cycles for each region. For example, delta for the USA is €35.

Step 4: We maximize profits using dynamic and myopic optimization – the difference between the two represents free-riding. We graph this difference on the right y-axis in Figure 9. The USA gains 51% in profits using dynamic rather than myopic optimization.

[18] As marketing activity varies by country, there is a potential concern for endogeneity in aggregate, cross-country analysis. By running separate models and optimizations per geographical regions, we are accounting for country-specific variations in marketing efforts. We thank an anonymous reviewer for this suggestion.

Peer Review Version

Step 5: Figure 9 shows the correspondence between delta, the excess of virtuous over vicious cycles, and the profit benefit achieved by dynamically taking advantage of this excess. The free-riding benefit mirrors the difference between virtuous and vicious cycles. [19]

Figure 9: Quantifying the "Free-Riding" Opportunity

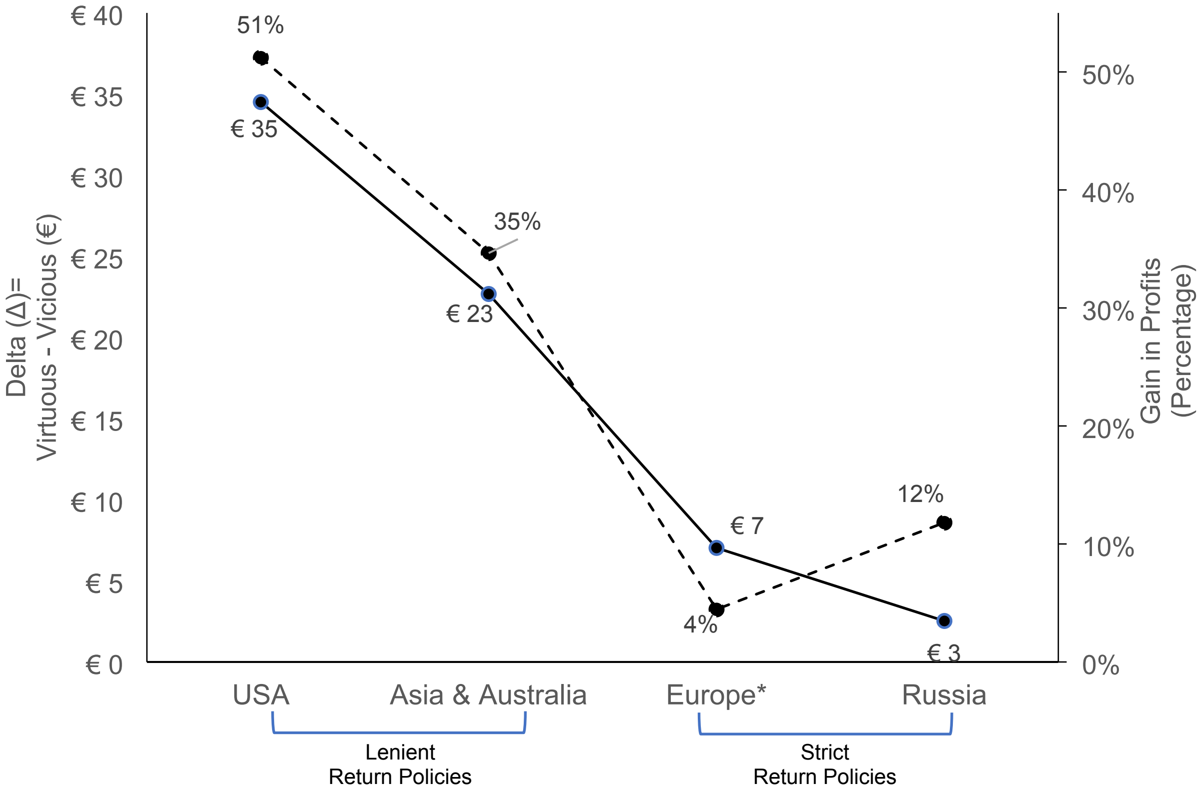

Notes:
* Europe includes only countries with stricter return regulations. The left y-axis is the excess of the virtuous cycle over the vicious cycle: $Delta = (Virtuous\ Cycle - Vicious\ Cycle) = (P_{it} \times OS_{it}) - (R_{it} \times RS_{it})$. The right vertical axis reports the "Gain in Profits" : $\frac{Profits\ (M1\ Dynamic) - Profits\ (M1\ Myopic)}{Profits\ (M1\ Myopic)} \times 100$.

Figure 9 shows lenient policies are associated with larger gaps between virtuous and vicious cycles and larger profit gains from dynamic optimization that exploits free-riding. We have: delta => gain from dynamic optimization => more free riding. To further investigate,

[19] We perform a sensitivity analysis where we compare our best model (M1/Dynamic) with another benchmark (M3/Dynamic). The pattern shown in Figure 9 remains consistent across this comparison.

we created Figure 7 by region (Web Appendix E, Figure W.E1). We found a pulsing pattern in Return Stock for lenient policy regions and none in strict policy regions.

These results suggest a large, positive difference between virtuous and vicious cycles creates free-riding opportunities. This difference is related to the return policy. Lenient return policies amplify the difference, enabling firms to capitalize on returns – free-ride them – to increase profits. Conversely, strict policies shrink this advantage, limiting the potential to fully leverage returns. The patterns in Figure 9 are striking. It is based on a "natural experiment" due to country differences, so we cannot unequivocally claim causality. However, it encourages future randomized field tests.

Discussion and Implications

Summary: The conventional wisdom is that returns are "bad" – a vicious cycle – but the intriguing finding is current returns are associated with higher sales in the long run – a virtuous cycle. The thesis of this paper is, if marketing is specified optimally, the firm can create and translate the virtuous cycle into higher profits, that is, "free-ride" returns.

We investigated whether a dynamic descriptive model combined with dynamic optimization can increase profits by free-riding. The answer is yes. The evidence was: (1) Optimal marketing in the presence of a virtuous cycle took on a pulsing pattern. "Return Stock", representing consumers' memory of returns, pulsed during the low-marketing periods, amplifying the virtuous cycle. (2) Dynamic optimization generated pulsing, whereas myopic did not, suggesting dynamic optimization *plus* a dynamic statistical model are needed to achieve free-riding. (3) Dynamic optimization refrained from over-investing in free-shipping promotions because they generated too many returns – too large a vicious cycle. (4) Lenient return policies were associated with larger differences between virtuous and vicious cycles,

Peer Review Version

and with larger profit differences between dynamic and myopic optimization. This suggests lenient policies produce more free-riding.

Our work fits the Decision Support System (DSS) paradigm. A DSS incorporates four ingredients (Little 1979; Lilien et al 1996). *Model:* This is the conceptualization of the key phenomena that need study. In our case, they are virtuous and vicious cycles. *Statistics:* This is the quantification of the model, the four-equation model we developed and our estimation of it. *Data:* These are the 1500 customers we analyzed. *Optimization:* This is the stochastic dynamic program we employed to generate an optimal free-riding strategy.

The DSS produces policies and insights. Tables 5 and 6 display policies for optimal marketing. Insights include: (1) Marketers may be able to free-ride the virtuous cycle if they develop marketing plans that generate a large virtuous cycle without generating too large a vicious cycle. (2) The goal is *not* zero returns, but a marketing strategy that exploits the virtuous cycle while accounting for the vicious cycle. (3) The right marketing policy may require re-allocation of spend among marketing activities, not necessarily decreased spending. (4) The opportunity to free-ride depends on the firm's return policy. Strict policies provide less opportunity to free-ride – they do not generate a strong enough virtuous cycle. (5) The build-up and decay of consumer memory quantified by "Return Stock" generates virtuous and vicious cycles. (6) Free-riding requires dynamic statistical *and* optimization models. This generates superior profits than when the empirical model, optimization, or both are not dynamic. Free-riding is a dynamic strategy – it requires dynamic measurement and optimization.

Our work contributes in three ways. First, it integrates dynamic optimization with empirical research that identified virtuous and vicious cycles. Second, consistent with the DSS perspective, it provides insights, as noted above. For example, our finding regarding the value of lenient return policies is provocative because many managers naturally shy away from

Peer Review Version

making it too easy to return. Third, also consistent with DSS, we showed how to implement the free-riding strategy, particularly the requirement to include dynamics both in the statistical model and in optimization.

Research Implications: First, product returns need to be included in transaction-data models of consumer purchase. They are associated with an immediate and long-term impact on sales. Second, the long-term impact requires that customer-targeted marketing maximize customer lifetime value (CLV). Third, while there is strong support for the virtuous cycle, it is less obvious than the vicious cycle, hence future research needs to measure it in different contexts – for example multichannel retailing environments. Additionally, the exact mechanism that creates the virtuous cycle needs to be ironed out. Bower et al. (2012) and Wood (2001) make a promising start regarding perceived fairness. However, there are several possible mechanisms and more need to be investigated. We shortly provide insights on these mechanisms using an exploratory online experiment.

Fourth, dynamic statistical modeling coupled with dynamic optimization is the sweet spot. Our results underscore how acute the differences are if one relaxes either on statistical or optimization dynamics. Fifth research needs to investigate the role of competition. Does competition learn the focal firm is pulsing, hence market out of phase in the belief they can over-ride the focal firm's free-riding? What is the marketing and profit equilibrium when multiple firms can free-ride? Sixth, future research should consider customer heterogeneity. Certain customers may be susceptible to the virtuous cycle, while others may to generate more vicious cycle. Seventh, we compiled our analysis sample to focus on every-year buyers. One of the benefits of this was to mitigate survivor bias by excluding "churners". However, controlling churn is a significant issue, hence future work should examine the relationship between vicious/virtuous cycles and customer churn.

Peer Review Version

Our analysis, and much of the empirical work on returns, employs observational data. To explore this further, we conducted two investigations detailed in Web Appendix F. First was propensity scoring matching (PSM). The treated group was customers who purchased in a given month and returned product. The control group was customers who purchased in that month but did not make a return. We repeated the analysis for several different time periods using different specifications of the PSM. Our findings consistently revealed a positive average treatment effect (ATT). Its magnitude was comparable to what one would infer from our estimated model in Table 4.

Secondly we conducted a pre-registered, randomized online experiment. We manipulated the returns policy – lenient *versus* strict. We expected the lenient policy to generate more returns, and it did. Our manipulation thus exogenously created high and low returns groups. We found the lenient group's self-reported purchase intentions were larger than those of the strict group, supporting the virtuous cycle. In addition, the test generated insights on the mechanisms that create the virtuous cycle. Mediation analysis suggested perceived fairness, relationship quality and lower perceived risk of purchasing the wrong product were key mechanisms translating the lenient policy to larger future purchase intentions.

These analyses support our findings of a positive relationship between returns and purchases – a virtuous cycle. However, they do not *prove* causality in the real world. The PSM is still subject to unobserved causal factors; the online experiment is subject to measurement specifics and questions of external validity. We strongly recommend *field experiments* to provide convincing evidence of virtuous and vicious cycles.

Managerial Implications: The good news for managers is that returns are not the enemy; the enemy is the opportunity loss of not managing them effectively. Unfortunately, there is no simple prescription. We can say marketing is the fulcrum but cannot generalize how

Peer Review Version

much to spend on which types of marketing. The free-riding strategy makes sense, but it takes a carefully optimized marketing schedule. With so much at stake, we suggest managers replicate our statistical model, combine it with dynamic optimization, and test its implications. Our methods, while nontrivial, are not out of the reach of well-trained data scientists.

Our DSS viewpoint guides managers on how to utilize the raw numbers the firm collects from its customers. The mantra is: model, statistics, data, and optimization. The model is "a preconceived idea of how the world works and therefore of what is interesting and worthwhile in the data." (Little p. 10). Our Figure 1 framework drives our DSS. It points to the key phenomena – virtuous and vicious cycles – the manager needs to balance. Statistics means specifying and estimating a set of equations ((1)-(6)) that follow from the framework. The optimization is any effort to improve the firm's current strategy. The statistical model provides parameter values required to run the optimization, and on its own provides important insights noted earlier. While the statistical model is an important component of the DSS, it is based on observational data – hence it provides valuable descriptive insights that support decision-making, but we cannot guarantee the results are causal.

The goal of a DSS is to *improve* managerial decisions. It thus suggests we compare "optimal" marketing policy to that currently employed by the focal firm. Table 7 does this. The optimal columns are from the "Dynamic Opt." allocation columns in Table 6. The actual is from records the retailer shared with us. The retailer appeared to be underspending on marketing. Perhaps the retailer was unaware of the power of marketing. Interestingly however, the retailer allocated funds roughly consistent with our DSS-generated recommendations – 78% to brand building efforts such as retailer brand and manufacturer brand advertising; 22% to the more promotional price discounts and free shipping. We would recommend taking funds away from discounting and putting them into free shipping, *although not overdo it*. Insights such as this are the hallmark of the DSS paradigm.

Peer Review Version

Table 7: Marketing Spend Allocation: Actual vs. Optimal: Average per Customer Per Week.

Type of Marketing Activity		**Actual**[a] **€ Spent**	**Actual**[a] **Proportion (%)**	**Optimal € Spent**	**Optimal Proportion (%)**
Non Brand-Building	Marketing Discounts	0.13	11.70	0.22	2.40
	Free Shipping	0.11	9.85	2.13	22.95
Brand - Building	Retailer Advertising	0.43	39.82	3.50	37.67
	Brand Advertising	0.42	38.63	3.44	36.99
Total		1.09	100.00	9.29	100.00

[a] To ensure consistency with the average optimal marketing, we calculated the average amount spent on various marketing activities during weeks 4-25 in the observed sample.

By comparing results at the country/region level, we show that a strict returns policy generates less virtuous cycle, thus less free-riding. In countries where the firm adopts a lenient policy, virtuous minus vicious is larger, providing the opportunity to free-ride returns. The gains from pursuing this opportunity are substantial (Figure 9). The key "takeaway" is a point of view: rather than seeing returns as something to be beaten down with strict policies that decrease returns but alienate consumers, the firm should consider lenient policies that create virtuous cycles and increase profits in the long run.

In summary, we have amplified previous findings in marketing to show returns can be managed to increase profits. The lever is marketing; the key phenomena are the dynamics of returns-induced virtuous and vicious cycles. Incorporating dynamics in both statistical models and optimization yields the most effective means for marketing to exploit the virtuous cycle and increase profits. Indeed, we hope our work stimulates researchers to continue the virtuous cycle of building on previous work to replicate, validate, and extend our results.

Peer Review Version

References

Ader, Jacob, Chai, Joyce, Adhi, Praveen, Touse, Sarah, and Hannah Yankelevich (2021). "Returning to Order: Improving Returns Management for Apparel Companies." *McKinsey Quarterly*, May.

Anderson, Eric T., Hansen, Kristan, and Simester, Duncan (2009). "The Option Value of Returns: Theory and Empirical Evidence." *Marketing Science*, 28(3), 405-423.

Ansari, Asim, Carl Mela, and Scott A. Neslin (2008), "Customer Channel Migration," *Journal of Marketing Research*, 45(1), 60-76.

Bower, Anthony B., and Maxham, James G. (2012). "Return Shipping Policies of Online Retailers: Normative Assumptions and the Long-Term Consequences of Fee and Free Returns." *Journal of Marketing*, 76(5), 110–124.

Blattberg, Robert C., and Neslin, Scott A. (1990). "Sales Promotion: Concepts, Methods, and Strategies."

Broadbent, Simon (1984). "Modeling with Adstock." *Journal of the Market Research Society*, 26(4), 295-312.

Chen, Jialie (2021). "A Structural Model of Purchases, Returns and Return-Based Targeting Strategies." *Journal of the Academy of Marketing Science*.

Danaher, Peter J., Bonfrer, André, and Dhar, Sanjay (2008). "The Effect of Competitive Advertising Interference on Sales for Packaged Goods." *Journal of Marketing Research*, 45(2), 211-225.

Dinner, Isaac M., Heerde Van, Harald J., and Neslin, Scott A. (2014). "Driving Online and Offline Sales: The Cross-Channel Effects of Traditional, Online Display, and Paid Search Advertising." *Journal of Marketing Research*, 51(5), 527–545.

Dubé, Jean-Pierre, Günter J. Hitsch, and Peter E. Rossi (2010), "State Dependence and Alternative Explanations for Consumer Inertia," *RAND Journal of Economics*, 41(3), 417-445.

El Kihal, Siham and Edlira Shehu (2022). "It's Not Only What They Buy, but It's Also What They Keep: Linking Marketing Instruments to Product Returns." *Journal of Retailing*, 98(3), 558-571.

El Kihal, Siham, Tulim Erdem, Christian Schulze and Weiqing Zhang (2025). "Customer Return Rate Evolution." *International Journal of Research in Marketing*, forthcoming.

Eley, Jonathan (2022). "Asos Warns on Profits as Customers Send Back More Purchases." Financial Times, June 19th. Retrieved November 9, 2022, from https://www.ft.com/content/33bc6b46-d23f-4178-94bf-4e7fbfb6a84c

Engel, James F., and Roger D. Blackwell (1982), *Consumer Behavior* (4th Edition), Hinsdale IL: Dryden Press.

Engel, James F., Roger D. Blackwell and Paul W. Miniard (1995), *Consumer Behavior* Dryden Press series in Marketing.

Erdem, Tülin, Susumu Imai, and Michel P. Keane (2003), "Brand and Quantity Choice Dynamics Under Price Uncertainty," *Quantitative Marketing and Economics* 1, 5-64.

Erdem, Tülin and Michael Pl Keane (1996), "Decision-making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science* 15(1), 1-20.

Fader, Peter S., James M. Lattin, and John DC Little (1992). "Estimating Nonlinear Parameters in the Multinomial Logit Model." *Marketing Science*, 11(4), 372-385.

Gonul, Füsun F. and Kannan Srinivasan (1996). "Estimating the Impact of Consumer Expectations of Coupons on Purchase Behavior: A Dynamic Structural Model." *Marketing Science*, 15 (3), 262-279.

Guide, Daniel R. Jr, Gilvan C. Souza, Luk Van Wassenhove, and Joseph D. Blackburn (2006). "Time Value of Commercial Product Returns." *Management Science*, 52(8), 1200-1214.

Janakiraman, Narayan, Holly A. Syrdal, and Ryan Freling (2016)."The Effect of Return Policy Leniency on Consumer Purchase and Return Decisions: A Meta-Analytic Review." *Journal of Retailing* 92(2), 226-235.

Khan, Romana, Michael Lewis, and Vishal Singh (2009). "Dynamic Customer Management and the Value of One-to-One Marketing." *Marketing Science*, 28(6), 1063-1079.

Keane, Michael P. (1997), "Modeling Heterogeneity and State Dependence in Consumer Choice Behavior," *Journal of Business and Economic Statistics*, 15(3), 310-327.

Kliegl, Oliver, and Karl-Heinz T. Bäuml (2021). "When Retrieval Practice Promotes New Learning–The Critical Role of Study Material." *Journal of Memory and Language* 120 (2021), 104253.

Lahiri, Kajal, and Peter Schmidt (1978). "On the Estimation of Triangular Structural Systems." *Econometrica: Journal of the Econometric Society*, 1217-1221.

Lilien, Gary L., Philip Kotler, and K. Sridhar Moorthy (1992). *Marketing Models*. Englewood Cliffs, New Jersey: Prentice Hall.

Little, John D.C. (1979). "Decision Support Systems for Marketing Managers." *Journal of Marketing*, 43(3), 9-26.

Mark, Tanya, Katherine N. Lemon, Mark Vandenbosch, Jan Bulla, and Anonello Maruotti (2013), "Capturing the Evolution of Customer-Firm Relationships: How Customers Become More (or Less) Valuable Over Time," *Journal of Retailing*, 89(3), 231-245.

Mulhern, Francis J., and Robert P. Leone (1991). "Implicit Price Bundling of Retail Products: A Multiproduct Approach to Maximizing Store Profitability." *Journal of Marketing*, 55(4), 63-76.

Neslin, Scott A., Thomas P. Novak, Kenneth R. Baker, and Donna L. Hoffman (2009). "An Optimal Contact Model for Maximizing Online Panel Response Rates." *Management Science*, 55(5), 727-737.

Ofek, Elie, Zsolt Katona, and Miklos Sarvary (2010). "'Bricks & Clicks': The Impact of Product Returns on the Strategies of Multi-Channel Retailers." *Marketing Science*, 30(1), 42-60.

Petersen, J. Andrew and V. Kumar (2009). "Are Product Returns a Necessary Evil? Antecedents and Consequences." *Journal of Marketing*, 73 (May 2009), 35-51.

Petersen, J. Andrew and V. Kumar (2015). "Perceived Risk, Product Returns and Optimal Resource Allocation: Evidence From a Field Experiment." *Journal of Marketing*, 73(3), 35-51.

Ratchford, Brian, Gonca Soysal, Alejandro Zentner, and Dinesh K. Gauri (2022). "Online and Offline Retailing: What We Know and Directions for Future Research." *Journal of Retailing,* 98, 152-177.

Roodman, David (2011). "Fitting Fully Observed Recursive Mixed-Process Models With CMP." *The Stata Journal*, 11(2), 159-206.

Sahay, B. S., Samir K. Srivastava, and Rajiv K. Srivastava (2006). "Managing Product Returns for Reverse Logistics." *International Journal of Physical Distribution & Logistics Management*, 36(7), 524-546.

Seetharaman, P. B. (2004), "Modeling Multiple Sources of State Dependence in Random Utility Models: A Distributed Lag Approach," *Marketing Science*, 23(2), 263-271.

For Peer Review Version

Shehu, Edlira, Dominik Papies, and Scott A. Neslin (2020). "Free Shipping Promotions and Product Returns." *Journal of Marketing Research*, 57(4), 640-658.
Simonov, Andrey, Jean-Pierre Dubé, Günter Hitsch, and Peter Rossi (2020), "State-Dependent Demand Estimation with Initial Conditions Correction," *Journal of Marketing Research*, 57(5), 789-809.
Statista (2022a). "Fashion E-Commerce Worldwide." Statista Dossier on Fashion E-Commerce Worldwide. Article Number: DID-111688-1.
Statista (2022b). "Returns of Online Purchases by Category." Published by Alexander Kunst.
Stephen, Andrew T., and Jeff Galak (2012). "The Effects of Traditional and Social Earned Media on Sales: A Study of A Microlending Marketplace." *Journal of Marketing Research* 49(5), 624-639.
Venkatesan, Rajkumar and V. Kumar (2004). "A Customer Lifetime Value Framework for Customer Selection and Resource Allocation Strategy." *Journal of Marketing*. 68(4), 106-125.
Wood, S. L. (2001). "Remote Purchase Environments: The Influence of Return Policy Leniency on Two-stage Decision Processes." *Journal of Marketing Research*, *38*(2), 157-169.

Peer Review Version

Web Appendix

Free-Riding Product Returns to Drive Profits

Elisa Montaguti
Scott A. Neslin
Sara Valentini

Peer Review Version

Author Contact Information:

Elisa Montaguti, Full Professor of Marketing, University of Bologna, Via Capo di Lucca 34, 40126 Bologna, Italy. Email: elisa.montaguti@unibo.it

Scott A. Neslin, Albert Wesley Frey Professor of Marketing, Tuck School of Business, Dartmouth College, Hanover, NH 03755, USA. Email: scott.a.neslin@dartmouth.edu

Sara Valentini, (Corresponding Author), Associate Professor of Marketing, Bocconi University, Via Roentgen 1, 20136 Milano, Italy. Email: s.valentini@unibocconi.it

Web Appendix

These materials have been supplied by the authors to aid in the understanding of their paper.

The AMA is sharing these materials at the request of the authors.

Peer Review Version

Web Appendix A: Data Coverage, Sample Construction, and Additional Descriptives

Descriptive Statistics of CC_i in Equations (1)-(4)

Table W.A1: Average and Standard Deviation (SD) of CC_i variables in Equations (1)-(4)

Macro Category CC_i	**Variable**	**Description**	**Mean**	**SD**
Past Behavior: [These variables were calculated using the first 26 weeks of data for each customer. The first 26 weeks were excluded from subsequent analyses.]	Returns Value $Proportion_i$	The ratio customer *i*'s total value of returns to the total value of purchases in the first 26 weeks	7.3%	19.3%
	Total Purchase $Occasions_i$	Customer *i*'s total number of purchases in the first 26 weeks	1.14	1.94
	Average Order $Size_i$	Customer *i*'s average OS in the first 26 weeks	83.43	154.10
	$Multichannel_i$ (Yes=1, No=0)	Indicates whether customer i utilized multiple channels for purchases in the first 26 weeks	0.13	0.50
	$Tablet_i$	Total number of purchases made through a tablet by customer *i* in the first 26 weeks	0.16	0.72
	$Smartphone_i$	Total number of purchases made through a smartphone by customer *i* in the first 26 weeks	0.15	0.74
	$Desktop_i$	Total number of purchases made through a desktop by customer *i* in the first 26 weeks	0.83	1.62
Customer Descriptors	$Gender_i$ (Female=1, Other=0)	1 indicates the customer *i* is a female	67.4%	46.9%
	Age_i (in years)	Age of customer *i* in years at the time of the purchase	29.02	20.52
	Age $Missing_i$ (Yes=1, No=0)	1 indicates customer *i* did not to provide date of birth	25.4%	43.5%
	Country: $Italy_i$ (Yes=1, No=0)	Indicates the country from which customer *i* purchased	29.3%	45.5%
	Country: $Japan_i$ (Yes=1, No=0)		6.2%	24.2%
	Country: USA_i (Yes=1, No=0)		14.4%	35.1%
	Country: $Russia_i$ (Yes=1, No=0)		16.0%	36.7%
	Country: $Germany_i$ (Yes=1, No=0)		6.2%	24.1%

Peer Review Version

Data Coverage & Sample Construction

Table W.A2: Sales coverage of the sample

Category	Description	N of Customers	Total Revenues (€)	Revenue Share
One and done	Purchased only in the acquisition year; no purchases in any later year of the observation window.	12581	2,109,858.50	18%
Intermittent multi years buyers	Purchased in one to three but not all subsequent calendar years after acquisition (years need not be consecutive).	2374	2,466,226.90	21%
Every-year buyers (Selected Sample)	Purchased at least once in every calendar year of the observation window; this is the modeling/selected sample.	1498	7,275,491.00	61%
Total	All customers across the 2013–2015 acquisition cohorts observed over 2013–2017.	16453	11,851,576.40	100%

Table W.A3: Churn (last observed purchase), purchase and return rates by year

Panel A: Cohort 2013 (2013–2017): survival, observed churn, purchase and return rates

Year (t)	**Customers alive at start of year t (N)[A]**	**Purchase rate among customers alive at start of year t (%)[B]**	**Return rate among purchasers alive at start of year t (%)[C]**	**Customers whose last purchase is year t (N) (Churners)[D]**
2013 Acquired	4744	100%	28%	3604
2014 Survivors	1140	76%	41%	179
2015 Survivors	961	77%	45%	181
2016 Survivors	780	83%	47%	201
2017 Survivors	579	NA	48%	NA

Panel B: Cohort 2014 (2014–2017): survival, observed churn, purchase and return rates

Year (t)	**Customers alive at start of year t (N)[A]**	**Purchase rate among customers alive at start of year t (%)[B]**	**Return rate among purchasers alive at start of year t (%)[C]**	**Customers whose last purchase is year t (N) (Churners)[D]**
2014 Acquired	5545	100%	28%	4237
2015 Survivors	1308	77%	42%	303
2016 Survivors	1005	80%	47%	276
2017 Survivors	729	NA	45%	NA

Peer Review Version

Panel C: Cohort 2015 (2015–2017): survival, observed churn, purchase and return rates

Year (t)	**Customers alive at start of year t (N)[A]**	**Purchase rate among customers alive at start of year t (%)[B]**	**Return rate among purchasers alive at start of year t (%)[C]**	**Customers whose last purchase is year t (N) (Churners)[D]**
2015 Acquired	6164	100%	27%	4740
2016 Survivors	1424	84%	42%	465
2017 Survivors	959	NA	43%	NA

Notes to Table W.A3:

[A] Customers whose first observed purchase ≤ t and last observed purchase ≥ t. This means that in the cohort's acquisition year (e.g., 2013 for the 2013 cohort) the purchase rate is 100% by construction.
[B] Purchase rate = $purchasers_t$ / $alive_t$ In the first row it equals 100% because the cohort is defined as customers whose first observed purchase is in that year.
[C] Return rate = $returners_t$/$purchasers_t$
[D] Customers whose final observed purchase falls in year t (no subsequent purchases within the window)

NA The observation window ends in 2017; therefore 2017 is right-censored and we do not report

Peer Review Version

Table WA.4 — Year-to-year transition probabilities between purchase/return states (Cohort 2013)

Panel A: Transition Matrix t=2013-->t+1=2014				
State in year t	**# of Customers in state in year t (N)**	**State in year t+1**	**# of Customers in state in year t+1 (N)**	**% transitioning to t+1 state**
Purchase & Return (t)	1309	Purchase & Return (t+1)	207	15.81
Purchase & Return (t)	1309	Purchase & No Return (t+1)	109	8.33
Purchase & Return (t)	1309	No Purchase (t+1) (churn)	993	75.86
Purchase & No Return (t)	3435	Purchase & Return (t+1)	143	4.16
Purchase & No Return (t)	3435	Purchase & No Return (t+1)	403	11.73
Purchase & No Return (t)	3435	No Purchase (t+1) (churn)	2889	84.10

Panel B: Transition Matrix t=2014-->t+12015				
State in year t	**# of Customers in state in year t (N)**	**State in year t+1**	**# of Customers in state in year t+1 (N)**	**% transitioning to t+1 state**
Purchase & Return (t)	350	Purchase & Return (t+1)	188	53.71
Purchase & Return (t)	350	Purchase & No Return (t+1)	73	20.86
Purchase & Return (t)	350	No Purchase (t+1) (churn)	89	25.43
Purchase & No Return (t)	512	Purchase & Return (t+1)	82	16.02
Purchase & No Return (t)	512	Purchase & No Return (t+1)	230	44.92
Purchase & No Return (t)	512	No Purchase (t+1)	200	39.06
No Purchase (t)	3882	Purchase & Return (t+1)	65	1.67
No Purchase (t)	3882	Purchase & No Return (t+1)	105	2.7
No Purchase (t)	3882	No Purchase (t+1) (churn)	3712	95.62
E.g., 3882 = 993 + 2889 from Panel A; 512 = 109 + 403 from Panel A; 350 = 207 + 143 from Panel A				

Panel C: Transition Matrix t=2015-->t+1=2016				
State in year t	**# of Customers in state in year t (N)**	**State in year t+1**	**# of Customers in state in year t+1 (N)**	**% transitioning to t+1 state**
Purchase & Return (t)	335	Purchase & Return (t+1)	191	57.01
Purchase & Return (t)	335	Purchase & No Return (t+1)	63	18.81
Purchase & Return (t)	335	No Purchase (t+1) (churn)	81	24.18
Purchase & No Return (t)	408	Purchase & Return (t+1)	66	16.18
Purchase & No Return (t)	408	Purchase & No Return (t+1)	177	43.38
Purchase & No Return (t)	408	No Purchase (t+1)	165	40.44
No Purchase (t)	4001	Purchase & Return (t+1)	45	1.12
No Purchase (t)	4001	Purchase & No Return (t+1)	105	2.62
No Purchase (t)	4001	No Purchase (t+1) (churn)	3851	96.25

Panel D: Transition Matrix t=2016-->t+12017				
State in year t	**# of Customers in state in year t (N)**	**State in year t+1**	**# of Customers in state in year t+1 (N)**	**% transitioning to t+1 state**
Purchase & Return (t)	302	Purchase & Return (t+1)	181	59.93
Purchase & Return (t)	302	Purchase & No Return (t+1)	57	18.87
Purchase & Return (t)	302	No Purchase (t+1) (churn)	64	21.19
Purchase & No Return (t)	345	Purchase & Return (t+1)	52	15.07
Purchase & No Return (t)	345	Purchase & No Return (t+1)	156	45.22
Purchase & No Return (t)	345	No Purchase (t+1)	137	39.71
No Purchase (t)	4097	Purchase & Return (t+1)	45	1.1
No Purchase (t)	4097	Purchase & No Return (t+1)	88	2.15
No Purchase (t)	4097	No Purchase (t+1) (churn)	3964	96.75

Peer Review Version

Web Appendix B: Complete Results for Models M1

Table W.B1: Proposed Model (M1)– Full Product Returns Dynamics: Complete Results.

Macro Category	Variable	Purchase Incidence (P_{it})		Order Size (OS_{it})		Return Incidence (R_{it})		Return Size (RS_{it})	
		Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Marketing Non-Brand Building	Marketing Discounts (β_{MKD})	0.010	0.004	0.006	0.005	-0.024	0.012	-0.012	0.008
	Free Shipping (β_{FS})	0.143	0.019	0.110	0.023	0.092	0.048	0.028	0.028
Marketing Brand-Building	Retailer Advertising (β_{RET})	0.014	0.002	0.006	0.002	-0.002	0.005	0.000	0.003
	Brand Advertising (β_{BRD})	0.102	0.028	-0.030	0.032	0.013	0.066	0.007	0.045
Marketing LAG Non-Brand Building	Marketing Discounts Lag ($\beta_{MKD.LAG}$)	0.016	0.004	0.017	0.005	0.012	0.012	-0.001	0.008
	Free Shipping Lag($\beta_{FS.LAG}$)	-0.022	0.020	-0.021	0.025	0.001	0.053	0.032	0.028
Marketing LAG Brand-Building	Retailer Advertising Lag ($\beta_{RET.LAG}$)	-0.005	0.002	-0.004	0.002	0.005	0.005	0.001	0.003
	Brand Advertising Lag ($\beta_{BRD.LAG}$)	-0.064	0.029	-0.053	0.033	0.031	0.077	-0.098	0.067
Product Ret. Dynamics	Return Stock λ = .93 (γ)	0.226	0.007	0.118	0.015	0.249	0.020	0.090	0.009
State Dependence	Purchase Incidence Lag (ι^{P})	0.320	0.014						
	Order Size Lag (ι^{OS})			0.082	0.004				
	Return Size lag (ι^{RS})							0.737	0.025
Recursive Component	Order Size (τ^{R}, τ^{RS})					0.325	0.061	0.042	0.004
Past Behavior [a]	Returns Value Proportion (κ_1)	-0.142	0.068	-0.292	0.118	1.293	0.175	0.174	0.047
	Total Purchase Occasions (κ_2)	-0.085	0.126						
	Average Order Size (κ_3)			0.001	0.000				
	Multichannel (Yes=1, No=0) (κ_4)	-0.020	0.026	-0.008	0.095	0.051	0.038	0.040	0.023
	Tablet (Yes=1, No=0) (κ_5)	0.117	0.112	0.022	0.028	0.010	0.020	0.012	0.008
	Smartphone (Yes=1, No=0) (κ_6)	0.183	0.127	0.030	0.014	0.024	0.014	-0.012	0.005
	Desktop (Yes=1, No=0) (κ_7)	0.107	0.126	-0.020	0.008	-0.029	0.010	-0.001	0.004
	Gender (Female=1, Other=0) (κ_8)	-0.044	0.025	0.000	0.041	0.102	0.060	-0.007	0.021
	Age (in year) (κ_9)	0.000	0.001	-0.002	0.001	0.004	0.002	-0.001	0.001

Peer Review Version

Macro Category	Variable	Purchase Incidence (P_{it})		Order Size (OS_{it})		Return Incidence (R_{it})		Return Size (RS_{it})	
		Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Observed Descriptors (CC_i)	Age Missing (Yes=1, No=0) (κ_{10})	-0.072	0.057	-0.096	0.074	0.055	0.099	-0.023	0.036
	Italy (Yes=1, No=0) (κ_{11})	-0.030	0.032	-0.321	0.040	0.312	0.072	-0.078	0.025
	Japan (Yes=1, No=0) (κ_{12})	0.218	0.044	0.146	0.142	-0.252	0.131	-0.051	0.037
	USA (Yes=1, No=0) (κ_{13})	0.030	0.045	0.195	0.060	0.006	0.110	0.029	0.032
	Russia (Yes=1, No=0) (κ_{14})	0.117	0.046	-0.124	0.078	-1.030	0.110	-0.133	0.049
	Germany (Yes=1, No=0) (κ_{15})	-0.013	0.056	0.086	0.076	0.274	0.096	0.076	0.038
Time Trend Dummies (T_t)	January (ν_1)	-0.096	0.020	0.028	0.024	-0.150	0.050	-0.014	0.031
	February (ν_2)	-0.005	0.020	0.068	0.023	0.030	0.049	0.007	0.034
	March (ν_3)	0.135	0.019	0.040	0.022	0.045	0.049	0.032	0.030
	April (ν_4)	-0.081	0.019	-0.021	0.023	-0.092	0.051	-0.038	0.031
	May (ν_5)	-0.053	0.021	0.058	0.023	0.028	0.053	0.023	0.030
	June (ν_6)	0.017	0.020	0.007	0.023	0.010	0.051	-0.004	0.029
	July (ν_7)	-0.137	0.017	-0.040	0.022	-0.025	0.046	-0.012	0.029
	September (ν_8)	0.112	0.017	0.066	0.022	-0.012	0.046	0.016	0.029
	October (ν_9)	-0.025	0.017	0.086	0.023	-0.031	0.049	0.085	0.028
	November (ν_{10})	0.236	0.018	0.212	0.024	0.036	0.049	0.075	0.028
	December (ν_{11})	0.036	0.019	0.026	0.022	-0.066	0.048	0.005	0.029
	Year 2014 (ν_{12})	0.006	0.035	-0.126	0.044	0.165	0.083	-0.029	0.049
	Year 2015 (ν_{13})	0.032	0.036	-0.131	0.045	0.184	0.084	-0.027	0.047
	Year 2016 (ν_{14})	0.018	0.037	-0.171	0.047	0.242	0.087	-0.015	0.050
	Year 2017 (ν_{15})	-0.025	0.037	-0.197	0.048	0.256	0.088	-0.002	0.050
	Cohort 2014 (ν_{16})	-0.081	0.029	-0.090	0.086	0.094	0.069	0.026	0.026
	Cohort 2015 (ν_{17})	-0.177	0.030	-0.152	0.039	0.084	0.079	-0.009	0.023
	Intercept (α)	-1.394	0.059	4.198	0.125	-2.905	0.324	0.099	0.133

[a]We created these variables using the first 26 weeks we observe for customer i. Therefore, we burned the first 26 weeks of observations for each customer.
[b]We refer to the country of residence of each customer *i* in our sample. We also have information on the country of purchase but in 99% of cases, the two variables coincide.
[c]We Order Size and Return Size variables are expressed as the natural logarithm.
General Notes: Number of observations= 252148, LL = -126985.28, Wald χ^2 = 48343.84, $p < .001$

8

Peer Review Version

Table W.B2: Comparing Models Depending on How Much They Include Returns Dynamics.

	M3: No-Returns Dynamics, (Return Stock and Lagged Returns Excluded)	**M2: Limited Returns Dynamics (Return Stock λ = 0)**	**M1: Proposed Model Full Product Returns Dynamics**
LL	-126985	-126762	-124990
N	252148	252148	252148
AIC	254126	253820	250307
BIC	256135	255365	252008

Complete estimates for benchmark models (M2 and M3) are available upon request.

Peer Review Version

Web Appendix C: Optimization Details

We identify optimal weekly marketing allocations by maximizing expected profits using the Generalized Reduced Gradient (GRG) algorithm in Frontline Solver. The optimization is based on customer-level predictions of purchase incidence, order size, return incidence, and return size from Equations (1) – (4), including random effects. To recover individual-level random effects, we simulate 500 draws from the four-element multivariate normal distribution estimated from the joint system of equations model. Each draw represents a potential realization of the unobserved customer-specific effect across the four outcome equations. For each customer, we assign weights to these draws based on the likelihood of observing that customer's actual behavior under each draw. The final estimate of the random effect vector is the weighted average across the 500 draws. This yields a four-dimensional customer-specific vector that captures unobserved heterogeneity.

We assess the predictive accuracy of our model using observed outcomes in holdout data. The figure reports actual purchase and return behavior across deciles ranked by model predictions. The consistent decline across deciles confirms that the model successfully identifies customers with higher expected outcomes. (see Figure W.C1).

Figure W.C1: Out-of-Sample Lift Charts– Predicted vs. Observed Purchase Incidence, Order Size, Return Incidence, Return Size

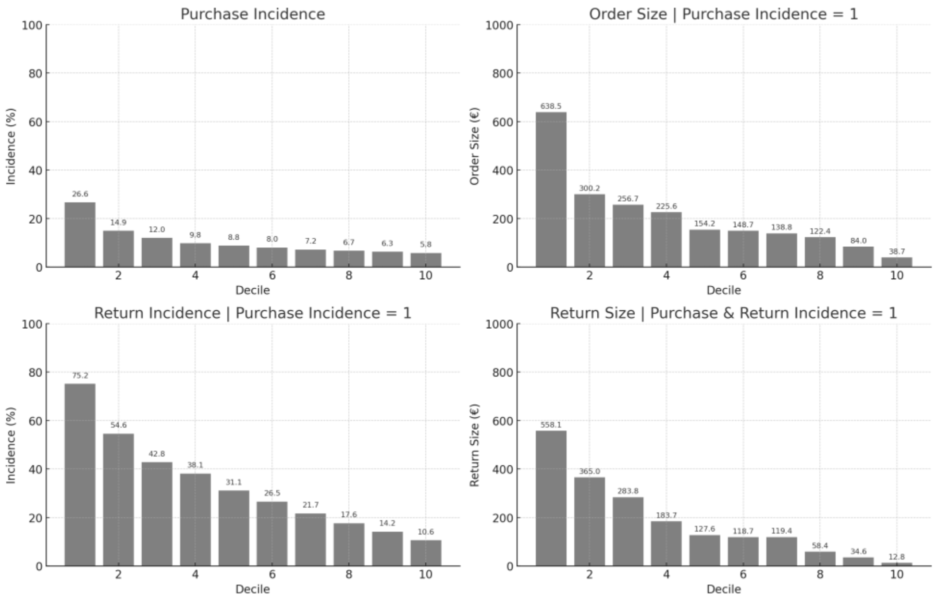

The objective function for the Generalized Reduced Gradient (GRG) algorithm is the profit function stipulated in the manuscript. We use the parameter estimates for Equations 1-6 to compute it. The decision variables are how much marketing to use each week. To address the risk of local optima, we initialize the optimization with multiple starting points for each customer. Given computational constraints, we apply this procedure to a random sample of 102 customers and optimize over 25 weeks per customer.

At each iteration, GRG updates a 4-dimensional vector Sjt, representing the number of days marketing instrument j is active in week t. For each week, we compute purchase probabilities using Equation 1, adjusting for state dependence. These probabilities are then converted into binary purchase outcomes via a random draw from a uniform (0,1) distribution. If a purchase is predicted, we compute order size using Equation 2. Similarly, return incidence is predicted via Equation 3, and return size is computed via Equation 4, conditional on a return.

Web Appendix D: Quantifying the Free-Riding Opportunity

To quantify the opportunity for free-riding returns, we compute the difference between the benefits and costs generated by $ReturnStock_{it}$. The benefits are what we call the virtuous cycle; the costs are the vicious cycle This difference (Δ) between them captures the net benefit of customer returns. While we find Δ always to be positive, it could be negative if the benefits of returns were less than the cost. Specifically, we compute:

$$\Delta_i = E[Revenues_{it}|ReturnStock_{it}] - E[ReturnCosts_{it}|ReturnStock_{it}|]$$

Expected revenue is computed as the product of purchase incidence and order size ($P_{it}\times OS_{it}$), while expected return costs are the product of return incidence and return size ($R_{it}\times RS_{it}$). We use the parameter estimates from our four-equation system (Equations 1–4) under the M1/Dynamic specification to compute Δ as follows: We simulate our estimated model assuming $ReturnStock_i$ is at its mean for each customer, then re-run the simulation assuming $ReturnStock_i$ increases by one standard deviation. This yields:

1. the *incremental revenue* from purchase incidence and order size as a result of a one SD increase in $ReturnStock_i$: $[P_{it}(+1SD) \times OS_{it}(+1SD)] - [P_{it}(Mean) \times OS_{it}(Mean)]$
2. the *incremental cost* from return incidence and return size as a result of a one SD increase in $ReturnStock_i$: $[R_{it}(+1SD) \times RS_{it}(+1SD)] - [R_{it}(Mean) \times RS_{it}(Mean)]$
3. We compute the net Δ_i as the difference between the two.

In the next section, we report results averaged at the country-group level to assess how return policy leniency moderates free-riding opportunities.

Web Appendix E: Regional and Policy Heterogeneity

Geographical Regions and Return Policies

The retailer's return policies vary across countries along two dimensions – customer cost to return (free vs. fee) and maximum number of post-purchase days when returns are allowed.

Peer Review Version

Countries can be classified as having Strict or Lenient return policies based on these dimensions. USA, Asia, and Australia have lenient policies; Russia and a subset of European countries have strict policies. We classify Europe other than the subset as in-between, i.e., "Mixed." We developed five groupings, which we call regions (Table W.E1).

Table W.E1: Geographical Regions and Return Policies

Region	Country	Maximum # Post-Purchase Days When Returns Allowed	Customer Cost for Returns	Classification	Region Percentage in Sample
Russia	Russian Federation	≤ 30	Fee	Strict	16.0
EU Countries with Strict Policy	Denmark, Croatia, Greece, Romania, Spain, The Netherlands, UK	≤ 14-15	Free	Strict	12.4
USA	USA	≤ 60	Free	Lenient	14.4
Asia & Australia	China, South Korea, Hong Kong Japan, Australia	≤ 30 ≤ 60	Free Free	Lenient	12.0
Europe	Germany, Italy, France, Austria, Belgium, Ireland Denmark, Croatia, Greece, Romania, Spain, The Netherlands, UK	≤ 100 ≤ 14-15	Free Free	Mixed	51.9
World	> 45 countries represented	From ≤ 14 to ≤ 100	Free and Fee	Mixed	100.0

We estimated M1 separately for each region; results are shown in Table W.E2. A region-specific grid search for λ, following the procedure used in Table 4, yielded consistent estimates across regions (λ = 0.91 for the USA and Europe; λ = 0.93 for Asia & Australia and Russia), closely aligning with the global model. While Return Stock coefficients in the purchase incidence equation are similar, the strength of the virtuous cycle differs across regions. This is due to three factors detailed in the manuscript (i) variation in average Return Stock levels — as shown in Figure W.E2, return incidence given purchase is much higher in lenient-policy countries (e.g., 39% in the USA) than in strict ones (e.g., 7% in Russia); (ii) differences in how Return Stock affects purchase vs. return — e.g., in the USA, its effect on purchase incidence (0.257) exceeds its effect on return incidence (0.187), unlike in Russia (0.246 vs. 0.232), and (iii) differing return policies. These factors explain the regional heterogeneity in outcomes, as illustrated in Figure 9.

For Peer Review Version

Table W.E2: Proposed Model (M1): Results Compared Across Regions.

Variable	Model	Regions with Lenient Return Policy		Regions with Strict Return Policy	
		USA	Asia & Aus.	Subset–EU	Russia
Marketing Discounts (β_{MKD})	Equation (1): P_{it}	-0.009 (0.010)	0.019 (0.009)	0.033 (0.019)	0.021 (0.011)
	Equation (2): OS_{it}	-0.014 (0.011)	0.017 (0.009)	0.001 (0.021)	0.031 (0.012)
	Equation (3): R_{it}	-0.025 (0.028)	-0.048 (0.027)	-0.016 (0.040)	-0.107 (0.043)
	Equation (4): RS_{it}	-0.018 (0.014)	0.005 (0.017)	-0.050 (0.024)	0.019 (0.026)
Free Shipping (β_{FS})	Equation (1): P_{it}	-0.006 (0.006)	0.225 (0.036)	0.008 (0.354)	0.181 (0.045)
	Equation (2): OS_{it}	0.055 (0.041)	0.065 (0.042)	0.039 (0.345)	0.106 (0.050)
	Equation (3): R_{it}	0.315 (0.103)	0.075 (0.035)	0.050 (0.423)	-0.028 (0.160)
	Equation (4): RS_{it}	-0.046 (0.047)	0.087 (0.072)	-0.011 (0.145)	-0.143 (0.144)
Retailer Advertising (β_{RET})	Equation (1): P_{it}	0.022 (0.005)	0.014 (0.003)	0.020 (0.006)	0.025 (0.005)
	Equation (2): OS_{it}	0.007 (0.006)	0.006 (0.005)	-0.007 (0.008)	0.020 (0.007)
	Equation (3): R_{it}	-0.020 (0.016)	-0.005 (0.011)	-0.017 (0.014)	-0.028 (0.019)
	Equation (4): RS_{it}	0.010 (0.008)	0.000 (0.007)	0.007 (0.009)	-0.005 (0.017)
Brand Advertising (β_{BRD})	Equation (1): P_{it}	0.106 (0.078)	0.158 (0.052)	0.049 (0.089)	0.237 (0.075)
	Equation (2): OS_{it}	-0.118 (0.102)	-0.061 (0.061)	0.162 (0.068)	-0.079 (0.066)
	Equation (3): R_{it}	-0.071 (0.198)	-0.002 (0.115)	0.366 (0.121)	-0.087 (0.271)
	Equation (4): RS_{it}	0.091 (0.081)	0.053 (0.111)	-0.218 (0.104)	0.383 (0.215)
Return Stock (γ)	Equation (1): P_{it}	0.257 (0.017)	0.240 (0.016)	0.218 (0.025)	0.246 (0.062)
	Equation (2): OS_{it}	0.103 (0.030)	0.062 (0.038)	0.061 (0.085)	0.115 (0.087)
	Equation (3): R_{it}	0.187 (0.054)	0.223 (0.062)	0.195 (0.078)	0.232 (0.152)
	Equation (4): RS_{it}	0.058 (0.017)	0.093 (0.028)	0.062 (0.021)	0.067 (0.104)
Decay λ (grid search)	All	0.910 (n.a.)	0.930 (n.a.)	0.910 (n.a.)	0.930 (n.a.)
Intercept (α)	Equation (1): P_{it}	-1.081 (0.110)	-1.275 (0.091)	-1.064 (0.133)	-2.810 (0.238)
	Equation (2): OS_{it}	4.383 (0.228)	4.719 (0.191)	4.269 (0.560)	5.044 (0.317)
	Equation (3): R_{it}	-2.493 (0.496)	-5.592 (0.559)	-2.321 (0.684)	-3.875 (1.003)
	Equation (4): RS_{it}	-0.587 (0.221)	-0.050 (0.336)	-0.411 (0.210)	0.257 (0.566)
Observations	–	36,192	69,524	28,470	40,326
Individuals	–	216	416	190	239
Log-Likelihood	–	-17,966.45	-35,946.07	-13,189.35	-18,780.16

Notes:

- Standard errors are reported in parentheses
- For space efficiency, estimates for the marketing-lag, state-dependence, and recursive components, as well as κ^P, κ^R, κ^{OS}, κ^{RS} (control variables) and ν^P, ν^R, ν^{OS}, ν^{RS} (time trend) parameters, are omitted but available upon request.

Peer Review Version

Figure W.E1: Dynamic Optimization: Average *ReturnStock$_{it}$* and Marketing Spend

Panel A: Regions with Lenient Return Policies[A]
Lenient Return Policies: USA

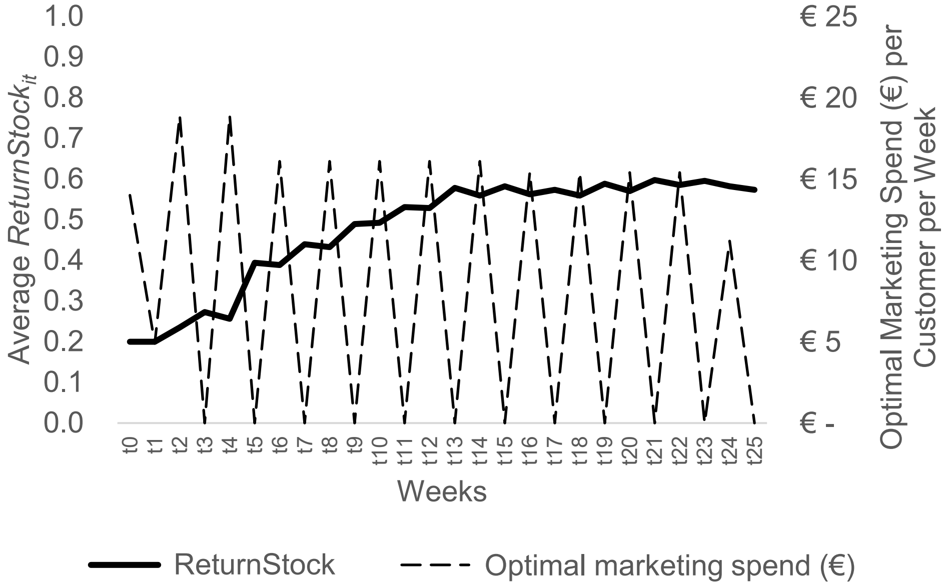

Lenient Return Policies: Asia & Australia

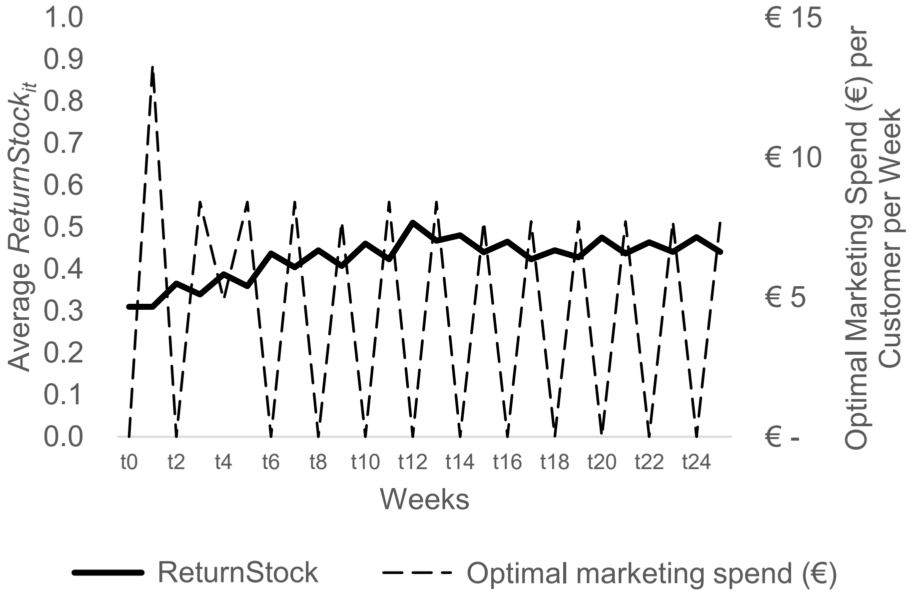

[A] A lenient return policy in our sample is characterized by free returns and a return period of 30 days or more.

Panel B: Regions with Strict Return Policies[B]
Strict Return Policies: Russia

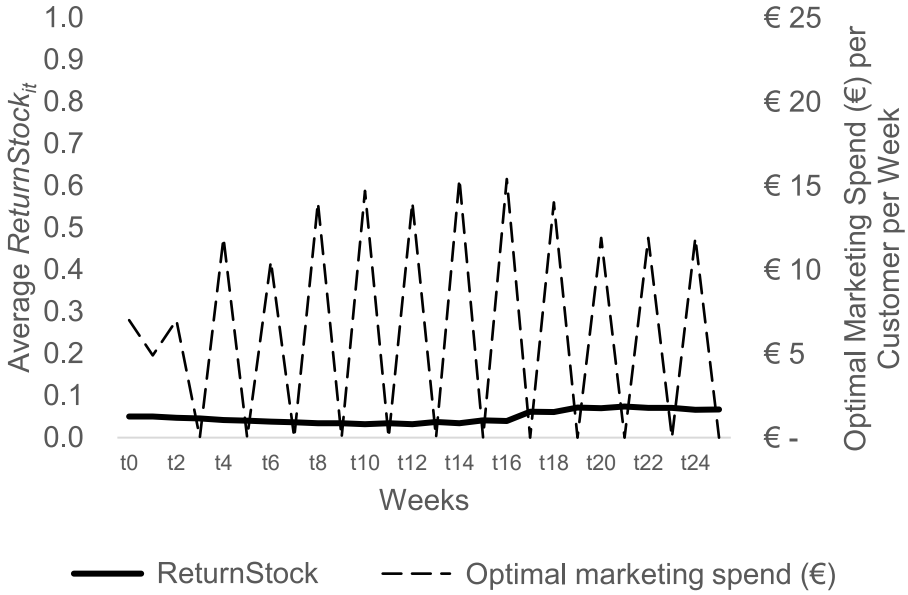

Strict Return Policies: European Countries with Strict Return Policies (e.g. UK, the Netherlands)[C]

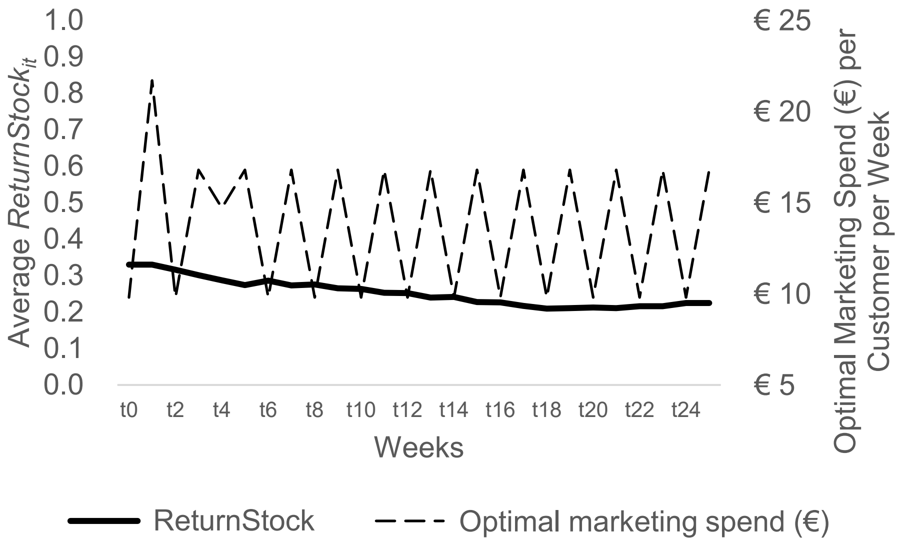

[B] A strict return policy in our sample is characterized by having to pay returns and/or a return period of 15 days or less.
[C] we included only European countries with a return period of 15 days or less under the category of "Strict Return Policies".

Figure W.E2: Average Return Incidence Conditional on Purchase, by Country (Grouped by Return Policy Leniency)

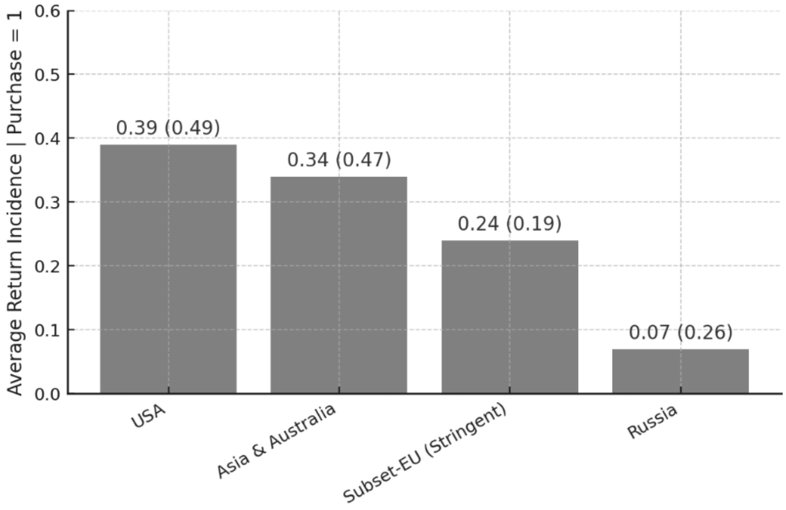

Note: Bars show the average return incidence conditional on purchase. Values in parentheses indicate standard deviation

er Review Version

Web Appendix F: Robustness & Exploratory Evidence: PSM and Online Scenario Experiment

The Impact of Returns on Purchase: Propensity Score Matching (PSM)

We employed PSM to assess the impact of returns on future purchases, i.e., the average treatment effect (ATT). The outcomes of interest are number of purchases and order size. For both outcomes, we run three different PSMs, using customer cohorts from 2014, 2015, and 2016. This yields a total of 2 × 3 = 6 PSMs. We use the first quarter of the year (Jan-Mar) to identify treated and control customers – this quarter is t=1. We assess outcomes for $t > 1$. Treated customers are those who purchased and returned during t =1; control customers are those who purchased in t = 1 but did not return. Each PSM estimates a logistic regression P(Return|X), then matches treated customers to controls whose logistic Return predictions as close as possible. The covariates (X) included:

- Customer Characteristics (age, gender, country)
- Marketing Exposure (current and lagged marketing interventions)

- Previous Purchase Behavior in *t* -1, and Previous Return Behavior in *t*-1[1]

We tested multiple PSM specifications (Nearest Neighbor with 1, 3, and 5 neighbors; calipers of 0.05, 0.1, 0.2; and Gaussian Kernel). Our final specification uses 5 neighbors with 0.2 caliper, achieving the best covariate balance based on mean bias, standardized bias (B%, i.e., the average standardized difference in covariates), and variance ratios (Austin, 2009; Caliendo & Kopeinig, 2008). The PSMs achieved good balance in all six analyses (2 outcomes × 3 years). On average, the mean standardized bias decreased from about 25% before matching to approximately 6% after matching, below the conventional threshold of 10%. The pseudo-R^2 for the logistic model dropped substantially (from ~0.13 to ~0.01), and likelihood ratio tests indicated no systematic differences in mean covariates between treated and matched control groups. Key covariates showed excellent balance, with standardized differences below 10% after matching. Variance ratios for continuous covariates generally fell within the recommended range [0.75, 1.33], confirming the quality of the matches across all periods. All PSM specifications showed significant positive ATTs of product returns on future purchases. Table W.F1 summarizes the results.

Table W.F1: PSM Results

Panel A: ATT for Number of Purchases in Subsequent Periods

Outcome: Total purchases in:	Treatment and Control Cohort					
	t=1, year==2014		t=1, year==2015		t=1, year==2016	
	ATT	t-stat	ATT	T=stat	ATT	t-stat
t+1	0.72	2.33	0.64	3.04	0.79	4.60
t+2	0.46	1.35	0.93	3.99	0.79	4.12
t+3	0.75	1.96	0.62	2.74	0.31	1.61
t+4	0.44	1.34	0.62	2.87	0.48	2.82
t+5	0.70	2.10	0.84	3.75	0.57	3.14
t+6	0.96	2.78	0.65	2.65	0.22	1.12
Average	0.67	1.98	0.72	3.17	0.52	2.90
Treatment (Sample Size)	76		196		374	
Control (Sample Size)	138		313		550	

[1] We also tested alternative PSM specifications: (a) excluding lagged return from the propensity model and (b) replacing it with Return Stock (t–1). These analyses yield ATT estimates that remain positive and statistically significant.

Peer Review Version

Panel B: ATT for Order Size per Purchase in Subsequent Periods

Outcome: Order size (net of return size) in:	Treatment and Control Cohort					
	t=1, year==2014		t=1, year==2015		t=1, year==2016	
	ATT	t-stat	ATT	t-stat	ATT	t-stat
t+1	€ 254.10	2.64	€ 200.94	3.32	€ 199.94	3.49
t+2	€ 247.27	2.23	€ 199.07	3.21	€ 236.16	3.41
t+3	€ 422.14	3.57	€ 277.73	3.71	€ 119.98	1.89
t+4	€ 177.71	1.65	€ 227.77	3.29	€ 99.25	1.59
t+5	€ 245.84	1.87	€ 174.78	2.63	€ 121.48	2.16
t+6	€ 269.11	3.26	€ 239.12	2.97	- € 61.67	-0.95
Average	**€ 269.36**	**2.54**	**€ 219.90**	**3.19**	**€ 119.19**	**1.93**
Treatment (Sample Size)	76		196		374	
Control (Sample Size)	138		313		550	

Notes to Table W.F1: Rows represent the quarters following the Return treatment. Columns represent the treatment and control groups for each cohort (2014, 2015, 2016). For example indicates customers who returned product they had purchased during *t*=1 purchased 0.72 additional times in *t*+1 compared to matched control customers purchased but did not return in *t*=1. ATT is calculated for each of six subsequent periods (t+1, t+2, ..., t+6), where t+1 corresponds to the first quarter after the treatment, t+2 to the second quarter, and so on.

The PSM results support our main findings (Table 4 in the manuscript). Both the number of future purchases (Panel A) and order sizes (Panel B) show consistently positive and significant relationships between product returns and subsequent purchase propensity.

To evaluate consistency between the estimated models in Table 4 of the manuscript and PSM results in Table W.F1, we calculated what the models would have suggested was the impact of returns on number of purchases. To do this we first simulated purchases for quarter *t*+1 assuming Return Stock at matched sample levels. We then ran a second simulation assuming no return experience (Return Stock = 0). All other variables were set to their means. The simulation predicted a 0.06 difference in weekly purchase probability between customers with (0.21) and without (0.15) return experience. When converted to quarterly frequency (×13 weeks). This translates to 0.77 additional purchases in quarter t+1, which aligns well with the PSM estimates (0.64-0.79 additional purchases reported in the first row of Table W.F1), confirming the robustness of effect size across different methodological approaches.

Peer Review Version

The Impact of Returns on Purchase: Online Experiment

Overview: Product returns are endogenous in observational data (customers self-select whether to return items). Therefore, we implemented a online experiment to further explore the impact of return experience on purchase intentions. Rather than directly forcing a return (which would be artificial and hence of limited external validity), we manipulated the strictness of the retailer's return policy in a between-subjects design. The goal was exogenously to create a high return group and a low return group. We pre-registered (https://aspredicted.org/p3mf-d9kc.pdf) the experiment and fielded it with 300 U.S.-based participants recruited via Amazon Mechanical Turk. Approximately 55% were female; 40% aged 26–40, 36% aged 18–25, and 19% aged 41–55; only 4% were aged 56 or older.

Procedure: Participants were randomly assigned to read one of two shopping scenarios featuring a fictional online apparel retailer ("Bon.Com") with either a lenient or a strict return policy. In the *lenient policy condition*, Bon.Com's return procedure was accommodating – it provided a direct "Returns" link on the website, included a prepaid return label in the package, allowed returns within 90 days, charged no return fees, and issued refunds immediately upon receiving the item. In contrast, the *strict policy condition* made the process more onerous – there was no direct Returns link (participants had to navigate through the website to find return instructions), returns were only accepted within 20 days, a $10 restocking fee was deducted from refunds, and refunds were delayed (issued 30–40 days after receiving the return). Both conditions were realistic, inspired by actual return policies. We expected the lenient condition to exogenously increase participants' likelihood of returning an item, enabling us to further explore the relationship between return experience and subsequent purchase.

After exposing respondents to one of the two returns policies, the survey informed respondents that a few months had gone by and asked them whether they would be likely to

Peer Review Version

purchase a light sweater from Bon.Com (5-point Likert scale). We then asked respondents to assume they had bought the sweater, but a friend told them it "doesn't suit you perfect" and that they began to second-guess their choice. We reminded them of the retailer's return policy and asked them if they were likely to return the sweater. We then asked respondents to agree or disagree with statements that reflected potential mediators between the firm's return policy and their decision of whether to return (e.g., "Bon.Com treats its customers fairly").

Manipulation Checks: Participants exposed to the lenient return policy reported significantly higher likelihood of returning the item than those in the strict condition ($M_{lenient}$ = 3.64 vs. $M_{stringent}$= 2.68; $t(296) = 6.63$, $p < .001$). They also reported higher perceived return experience ($M_{lenient}$ = 2.18 vs. $M_{stringent}$ = 1.79; $t(292) = 2.77$, $p = .006$).

Analysis: The lenient policy significantly increased purchase likelihood ($M_{lenient}$ = 3.70 vs. $M_{stringent}$ = 2.36; $t(285) = 10.83$, $p < .001$). This suggests that lenient policies lead to more future purchases. Since respondents randomly assigned to the lenient condition stated they were more likely to return, we interpret this as evidence of a positive relationship between returns and future purchase.

To examine the mechanisms by which returns lead to future purchase, we conducted mediation analyses between policy and purchase. As discussed in the literature review section of the manuscript, fairness, risk reduction, and relationship quality are psychological mechanisms through which return experiences might influence future purchase behavior. The items in Table W.G1 show these mediators.

Table W.G1: Overview of Measures Used in the Online Experiment

Construct	Example Item
Purchase Intention	How likely would you be to buy a new item from Bon.Com?
Return Intention	How likely would you be to return the item?
Return Experience	I have ample experience with returning items to Bon.Com.
Fairness	Bon.Com treats its customers fairly.
Relationship Strength	I have a strong relationship with Bon.Com.
Risk Reduction	Buying at Bon.Com reduces the risk I end up with the wrong product.
All items were measured on a 7-point scale.	

Peer Review Version

We conducted mediation analyses using the PROCESS macro developed by Hayes (2017, Model 6; 5,000 bootstrap samples). Results are in Table W.G2. Each model reveals significant indirect effects of the return policy manipulation (lenient vs. strict) on purchase likelihood. The strongest mediators are return experience and mechanisms such as fairness, relationship quality, and risk reduction.[2] This is consistent with theory – return policies that precipitate returns increase future purchase because of perceived fairness, strong relationship with the retailer, and decreased risk of purchasing the wrong product. We also ran the same mediation models using return intention as the dependent variable.

Table W.G2: Mediation Analysis Results

Mediation Path	Indirect Effect	Boot SE	95% CI (LLCI)	95% CI (ULCI)
Return Policy → Return Exp. → Purchase Int.	.06	.03	0.01	0.12
Return Policy → Fairness → Purchase Int.	.63	.10	0.43	0.84
Return Policy → Return Exp. → Fairness → Purchase Int.	.05	.02	0.01	0.09
Return Policy → Relationship → Purchase Int.	.27	.07	0.15	0.41
Return Policy → Return Exp. → Relationship → Purchase Int.	.09	.04	0.03	0.17
Return Policy → Risk Reduction → Purchase Int.	.65	.10	0.47	0.85
Return Policy → Return Exp. → Risk Reduction → Purchase Int.	.04	.01	0.01	0.07

Conclusion: The online experiment, PSM analysis, the four-equation model, and previous research discussed in the manuscript, support the idea that returns can increase future purchases. Potential reasons are perceptions of fairness, relationship quality, and risk reduction. That said, we acknowledge like all such experiments, (1) it is subject to specifics such as the ordering of questions – e.g., we asked purchase likelihood after providing return policy information but before we asked about returns. (2) the experiment does not provide perfect external validity, i.e., real world behavior. We see the experiment as a complement to our statistical results – neither is perfect, but together provide encouraging evidence for the virtuous and vicious cycles.

[2] Each mechanism (fairness, relationship quality, and risk reduction) was tested in a separate mediation model, rather than jointly in a single structural model. This approach enables us to isolate the indirect pathway through each construct, though it does not account for potential intercorrelations among mediators.

Peer Review Version

Web Appendix G: Addressing Marketing Endogeneity

The retailer's management informed us that marketing was not targeted at individual customers. However, marketing differed by country. This variation could be due to unobserved country-specific demand shocks, creating endogeneity in marketing variables. We do not think this is a problem for two reasons: (1) Country-specific models also reveal sensible virtuous and vicious cycles, as discussed in the manuscript. (2) The results of our aggregate model are stable after using Gaussian copulas to address endogeneity.

(1) We estimate our aggregate model at the country level with time dummy variables to control for country-specific and time-varying factors (see Table W.E2). This approach allows us to account for country-specific patterns in marketing decisions and control for time-varying unobservables within each country. It is noteworthy that virtuous and vicious cycles exist at the country level as well as in aggregate.

(2) While this mitigates endogeneity concerns with country-specific models, there still could be endogeneity in the aggregate model in Table 4. A standard remedy is to use instrumental variables (IV), but we do not have promising instruments. We thus turn to an instrument-free approach, namely Gaussian copulas (Park & Gupta 2012; Papies, Ebbes, & Van Heerde 2017). Following Park and Gupta (2012), we construct copula correction terms and include them in the model. The copulas are computed at the customer-week level, preserving both the time and customer dimensions of the data.

Table W.G1 shows results for the aggregate model with and without the copulas. Marketing coefficients are very similar in both cases, suggesting endogeneity is unlikely to be driving our findings. A few copulas attain marginal significance, but this just means the corresponding regressor was endogenous – the copula controls for this. Note the Return Stock variables are also stable with *versus* without copulas. We also replicated the analysis separately for each country, estimating five multi-equation models in total. Across these

For Peer Review Version

models, we tested 97 copulas, of which 7 reached significance at p = .05. The stability of our estimates with and without copulas suggests endogeneity is not a concern for the aggregate model shown in Table 4. Full results are available upon request.

Table W.G1: Table 4 Model (M1)– Marketing Estimates With/Without Copula Correction.

Variable	Purchase Incidence (Pit)		Order Size (OSit)		Return Incidence (Rit)		Return Size (RSit)	
	Without Copula	With Copula	Without Copula	With Copula	Without Copula	With Copula	Without Copula	With Copula
	Coef.	Coef.	Coef.	Coef.	Coef.	Coef.	Coef.	Coef.
Marketing Discounts (β_{MKD})	0.010***	0.011**	0.006	0.005	-0.024**	-0.069***	-0.012	-0.029
Free Shipping (β_{FS})	0.143***	0.139***	0.11***	0.084**	0.092*	0.113**	0.028	0.030
Retailer Advertising (β_{RET})	0.014***	0.016***	0.006***	0.010**	-0.002	-0.002	0.000	0.002
Brand Advertising (β_{BRD})	0.102***	0.097***	-0.03	-0.033	0.013**	0.006*	0.007	0.005
Marketing Discounts Lag ($\beta_{MKD.LAG}$)	0.016***	0.025***	0.017***	0.009*	0.012	-0.028	-0.001	-0.017
Free Shipping Lag($\beta_{FS.LAG}$)	-0.022	-0.031	-0.021	-0.126	0.001	0.293	0.032	0.021
Retailer Advertising Lag ($\beta_{RET.LAG}$)	-0.005***	-0.004	-0.004**	-0.002*	0.005	0.024	0.001	-0.004
Brand Advertising Lag ($\beta_{BRD.LAG}$)	-0.064***	-0.065***	-0.053	-0.046	0.031	0.010	-0.098	-0.090
COPULA Marketing Discounts		0.008		0.001		0.024**		0.013*
COPULA Free Shipping		0.069		0.005		0.100		0.080
COPULA Retailer Advertising		-0.002		-0.007		0.015		-0.004
COULA Marketing Discounts Lag		0.008**		0.013		0.022		0.009
COPULA Free Shipping Lag		0.070		0.026		-0.071		0.060
COPULA Retailer Advertising Lag		-0.003		-0.004		0.026*		0.007
Return Stock λ (γ)	0.226***	0.222***	0.118***	0.109***	0.249***	0.214***	0.09***	0.070***
State Dependence	0.32***	0.318***	0.082***	0.081***	na	na	0.737***	0.764***
Intercept (α)	-1.394***	0.000***	4.198***	4.448***	-2.905***	-2.769***	0.099	0.077

Notes to Table W.G1:

- $p < 0.01$, ** $p < 0.05$, *** $p < 0.001$
- κ^P, κ^R, κ^{OS}, κ^{RS} parameters for control variables and ν^P, ν^R, ν^{OS}, ν^{RS} for time dummy variables are not reported
- the *Copula Brand Advertising* term was dropped due to perfect collinearity with its associated regressor (*Brand Advertising*). This occurred only for brand advertising variables, where the transformation used to construct the Copula resulted in redundant information.

Peer Review Version

References

Austin, Peter C. (2009). "Balance Diagnostics for Comparing the Distribution of Baseline Covariates Between Treatment Groups in Propensity-Score Matched Samples." *Statistics in Medicine*, 28(25), 3083–3107.

Caliendo, Marco and Sabine Kopeinig (2008). "Some Practical Guidance for the Implementation of Propensity Score Matching." *Journal of Economic Surveys*, 22(1), 31–72.

Hayes, Andrew F. (2017). Introduction to Mediation, Moderation, and Conditional Process Analysis: A Regression-Based Approach (2nd ed.). New York: The Guilford Press.

Papies, Dominik, Peter Ebbes, and Harald J. van Heerde (2017). "Addressing Endogeneity in Marketing Models Using a Gaussian Copula Approach." *International Journal of Research in Marketing*, 34(1), 137–155.

Park, Sanghee and Sunil Gupta (2012). "Handling Endogenous Regressors by Joint Estimation Using Copulas." *Marketing Science*, 31(4), 567–586.

Peer Review Version