ATTENTION AND SALIENCY ON THE INTERNET: EVIDENCE FROM AN ONLINE RECOMMENDATION SYSTEM*

Christian HELMERS
Santa Clara University

Pramila KRISHNAN
University of Cambridge & CEPR

Manasa PATNAM†
CREST (ENSAE)

November 2015

ABSTRACT

Using high-frequency transaction-level data from an online retail store, we examine whether consumer choices on the internet are consistent with models of limited attention. We test whether consumers are more likely to buy products that receive a saliency shock when they are recommended by new products. To identify the saliency effect, we rely on i) the timing of new product arrivals, ii) the fact that new products are *per se* highly salient upon arrival, drawing more attention and iii) regional variation in the composition of recommendation sets. We find a sharp and robust 6% increase in the aggregate sales of existing products after they are recommended by a new product. To structurally disentangle the effect of saliency on a consumer's consideration and choice decision, we use data on individual transactions to estimate a probabilistic choice set model. We find that the saliency effect is driven largely by an expansion of consumers' consideration sets.

KEYWORDS: Limited attention, advertising, online markets.

JEL Classification: D22, M30, K11, O34

*We thank Romain Aeberhardt, Yann Bramoulle, Andrew Ching, Jishnu Das, Alessandro Iaria, Ali Hortacsu, Thierry Kamionka, Francis Kramarz, Laurent Linnmer, Oliver Linton, Marco Mariotti, Sendhil Mulainathan, Carlos Ponce, Matthew Shum, Anand Srivastava, Catherine Tucker, and seminar/conference participants at Aix-Marseille School of Economics (GREQAM), Cambridge-INET/Cemmap Workshop on Economic and Econometric Applications of Big Data, CREST, ENSAE-Berekely DataLead Conference, Georgetown University, the NBER Summer Institute, 2nd Empirical Microeconomics Workshop Banff, XX LACEA Annual Meeting, and the 13th Annual International Industrial Organization Conference for their useful comments and suggestions. We thank especially, Kim Montalibet and Elango Venkatachalapathy who provided excellent research support; their input was essential and invaluable. Financial support from the LABEX ECODEC (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047) and INET is gratefully acknowledged.

†Corresponding author: manasa.patnam@ensae.fr

1 Introduction

A standard simplifying assumption about consumer behaviour is that consumers consider all alternatives when making choices. However, exercising choice requires awareness of all available options which may be limited by search costs or by cognitive overload over the number and variety of products available to a consumer (Sims 2003; Caplin et al. 2011). Many e-commerce websites offer thousands of products for sale even within narrowly defined subcategories. For instance, Amazon offers more than 3,000 options for buying a television. Further narrowing of search in this category produces at least 190 options.[1]

A growing literature on the nature of consumer search on the internet has documented that despite the low physical search costs associated with internet browsing, there appears to be a prevalence of search frictions due to a scarcity of attention towards the large variety of available choices (Dinerstein et al. 2014). In the face of very large choice sets in these settings, it is plausible that consumers may try to simplify decisions by examining a smaller set of products (Masatlioglu et al. 2012); for example, Kim et al. (2010) report that consumers typically only search for 11% of all available options (for camcorders) on Amazon. Experimental evidence too, confirms that limits to attention impose a bottleneck on processing stimuli (Mozer and Sitton 1998).

In view of this, many online marketplaces frequently engage in tactics to attract user attention towards their products. Eliaz and Spiegler (2011) show that marketing devices, employed by firms, are likely to influence the set of products that a consumer chooses to consider, termed her *consideration set*. Hauser and Wernerfelt (1990) confirm this empirically, and find that consumer consideration sets respond to actions by firms that increase the visibility or salience of products.[2] One such tool, the product recommendation system,[3] through which a subset of products are selectively highlighted to users, is increasingly used by firms to attract attention and increase awareness towards products (Fleder and Hosanagar 2007). Yet, field evidence on the effectiveness of such recommendations, and indeed how they affect consumer choice, is limited. A small body of experimental studies in the laboratory find that subjects who receive recommendations for a product are more likely to select it relative to those who do not receive them (Senecal and Nantel 2004; Huang and Chen 2006). Kim et al. (2010) use aggregate product search data from Amazon and find in a policy simulation that a recommendation system, highlighting popular products, significantly affects their demand and lowers search costs.

In this paper, we examine whether consumer choice online is influenced by impersonal product recommendations. To investigate the question, we use data on shopping purchases from

[1] The applied search filters are: Home Entertainment, 50-59 inches, 1080p resolutions, flat screen.

[2] A situation where firms can influence the attention process by sending signals to consumers is commonly referred to as 'stimulus-driven' attention allocation in contrast to 'rational inattention', where (only) consumers choose and optimize resources for the allocation of their attention. The effect of visibility on consumer behavior is also known as the 'mere exposure effect' in the psychology literature (Zajonc 1968).

[3] Product recommendation systems typically use a 'collaborative-filtering' technology that chooses items for recommendations based on similarity measures between users and/or items (Schafer et al. 2007).

an exclusively online retail store, which offers a list of product recommendations for every product available on the website; these recommendations are not personalized and instead based largely on attribute similarity. We focus on two objectives: the first is to provide causal evidence on the the extent to which recommendations affect aggregate product sales, independent of product popularity or underlying characteristics. Our second objective is to examine whether the saliency effect generated by product recommendations limits consumer attention to a smaller consideration or evoked set. To do so, we estimate a reduced form model of consumer choice that allows us to determine the causal effect of product recommendations on sales. To establish that the effect is driven by the impact of saliency on consideration sets, we estimate a random utility model that incorporates the formation of a consideration set in the first stage of a consumers' decision making process and test whether saliency generated by product recommendations affects both the formation of this set and the consequent choice.

Our identification strategy focuses on product recommendations coming from the arrival of new products. Since new products are highly viewed upon arrival, we are able to focus on the aspect of product *saliency*, that is, the prominence of a product in a consumer's mind. Specifically, we analyze what happens to the sales of an existing product after it is recommended by new products. Since new products are highly salient, the recommended existing product receives a positive "saliency shock." We exploit both the timing of new product arrivals, which highlight a set of similar products already available on the website (the recommendation set), and regional variation between Europe and North America in recommendation sets to identify our saliency effect. Our double difference-in-difference strategy allows us to difference out both product-by-time and region-by-time unobservables. In addition we are able to employ new-product fixed effects to absorb any possible correlation between characteristics of new products and the products they recommend, ensuring that these saliency shocks are treated as exogenous with respect to the individual existing recommended products.

There is, by now, a large literature on the effect of salience measured via product popularity on sales (see for example, Sorensen (2007), Carare (2012), Cai et al. (2009) and Tucker and Zhang (2011)).[4] However, such evidence cannot rule out the possibility that the visibility of such goods actually imparts information. The effect of being on the first page or being popular also captures latent product quality or price-based relevance. Our approach differs from the existing literature as we focus on the effect of saliency shocks generated by recommendations from new product arrivals. The saliency of new products is useful because it generates a "spillover saliency" effect for existing products that are recommended by the new product. Our identification strategy, therefore, allows us to distinguish between pure saliency and information/popularity effects.

[4] Smith and Bryniolfsson (2001) examine the online purchase of books and find that being among the first few entries or on the first page on a search list is more important than being the first entry. Baye et al. (2009) show that on-screen placement of a link, based on its pricing, on an online search page is a central determinant of the number of clicks. There is also considerable evidence that salience and limited attention matter in other settings: Chetty et al. (2009) show that that consumers underreact to taxes that are not salient. DellaVigna (2009) provides a comprehensive review of the literature examining saliency and limited attention.

Our results indicate a sharp and robust 6% increase in sales of products after they receive a saliency shock. This saliency effect is short-lived, with the majority of the effect concentrated around the day that the product receives the saliency shock, diminishing rapidly thereafter. The effect completely disappears three days after the saliency shock has been received when the next batch of new products is launched on the platform. Such a pattern of effects, with a prominent spike in a product's sales on the day that they are recommended, is consistent with the "attention" based explanation that products receiving a saliency shock have been previously overlooked by consumers. The lack of persistent long-term effects on sales is in fact incompatible with selection-based explanations which would imply that recommended products would have experienced an increase in sales regardless of the recommendation.[5]

Further, we find that the saliency effect is larger for products recommended in smaller sets, suggesting that consumers pay attention to products that are more visible. We also find significant (positive) spillover effects. Products recommended by saliency-shock-affected products also see an increase in their sales on the day that new products are launched, but to a much lesser extent.

Although our reduced form results show that recommendations have a sizeable causal impact on sales, it is unclear how this effect is derived from a consumer's choice process. Products that gain saliency are likely to receive more *attention* by consumers, in this way increasing the probability that they are considered (Haan and Moraga-González 2011). Even then, it is possible that, independent of increased attention, consumers derive utility from saliency i.e, they have a direct preference for products recommended by new products.[6] Therefore, in order to understand how consumer demand is affected by recommendations, we estimate a random utility model that incorporates the formation of a consideration set in the first stage of a consumers' decision making process. Manzini and Mariotti (2014) provide theoretical underpinnings for the formulation of such types of stochastic consideration sets and show that the consideration probability can be interpreted as an attention parameter, indirectly measuring the degree of product awareness. Our consideration set model, based on Manski (1977), recognizes the choice process as sequential and allows for heterogenous consideration sets across individuals. In this way, we are able to distinguish between a consumer's consideration for the product due to saliency and her preference for saliency itself. Our model also avoids the Independence of Irrelevant Alternatives (IIA) assumption and accommodates the feature of choice frequency reversal due to the addition or elimination of other alternatives. Note that we do not however explicitly model the consumer search process, as is done, for example, in De los Santos et al. (2012) and Kim et al. (2010) and thus are unable to draw conclusions

[5] All our specifications include (panel) leads on the saliency shock variable, as well as their cumulative pre-saliency status, to mitigate the concern of anticipation effects. Our results consistently show insignificant effects for the included leads. In addition, our alternative identification strategy, which exploits regional variation in the composition of recommendation sets, relies only on a common trends assumption between a recommended and non-recommended product *for the same new product*. We show that this assumption is empirically validated in our data.

[6] For example, it is possible that all products linked to a new product are considered "trendy." To the extent that consumers have a direct preference over consuming products that are in line with the current trend, our proxy for saliency will enter a consumer's utility function.

about how recommendations may affect search.[7]

Using micro-data on shopping transactions we find that saliency has a strong, positive effect on the consideration (of existing products) but no further effect on choice, conditional on consideration. We estimate that the saliency effect is higher within sub-categories where only a few choices are considered. Based on our model estimates we present counterfactuals that compare how sales shares for products change when consumers have limited vs full information. Our results indicate that under the current recommendation system, where all products tend to be equally highlighted, popular products tend to suffer a loss in sales share when consumers have limited attention simply because they are not considered. Popular products would, however, stand to gain under limited attention (by up to 4% of sales share difference), if the website only recommended popular products, but this increases the concentration of sales towards popular products in the market. These results are consistent with a "segmentation effect" of improved within-platform search (through for example, recommendation systems) that shifts market participation in favor of some products against others, identified theoretically by Lewis and Wang (2013).

There is some existing empirical evidence to support models of limited attention through the formation of consideration sets. A large literature, mainly in marketing science, uses explicit structural or functional form restrictions to model the formation of consideration sets and its subsequent effect on choice (see for example Van Nierop et al. (2010); Chiang et al. 1998; Andrews and Srinivasan 1995; Roberts and Lattin 1991). A small but growing literature in economics uses exclusion restrictions to estimate a consumer demand model with limited attention. For example, Goeree (2008) uses advertising expenditure, proxied by media exposure, as an exogenous shifter that affects a consumer's consideration but not her utility. Draganska and Klapper (2011) do not impose such a restriction on advertising but treat it as fully exogenous and assume that the consumer's choice set is limited to only the set of brands that she is aware of. Kawaguchi et al. (2014) propose an alternative methodology that uses product availability as an exclusion restriction to test for attention.[8]

Our methodological approach is similar in spirit but we exploit a richer context, incorporating not just the availability of new products but also variation in product visibility, to identify and estimate the presence of consumer inattention. Additionally, in contrast to some of the existing literature, our paper does not require product saliency, a form of advertising, to be exclusive to the consideration process and allow it to affect consumer utility. Neither do we place any restrictions on the composition of the consideration sets or the process of choice formation. Finally, instead of relying on aggregate proxies, we use a direct measure of product visibility and salience shocks that vary frequently both across products and over time.

[7]We lack data on consumer specific search over different products to be able to explicitly feature search in our framework. In general, consumers may face heterogenous search costs that are fixed across all products (De los Santos et al. 2012) or vary by product (Kim et al. 2010), but our consideration stage specification is agnostic about this. In this sense, our results for the consideration effect can be viewed as a mixture of increasing awareness and lowering search costs.

[8]Conlon and Mortimer (2013) provide an application using stock-outs, but they do not take into account the endogeneity of product stock-outs.

Our research contributes to two important literatures. First, our paper contributes to the fast-growing literature on the economics of digitization that analyzes consumer behavior on the internet by offering evidence on the effect of online recommendation systems. To the best of our knowledge, our paper is the first to do so. Choosing amongst the entire range of products offered by e-commerce can often be challenging for any consumer shopping on the internet, and recommendation systems offer a possible way to alleviate this friction. We also contribute to the field of behavioural economics and the literature on the economics of attention by offering empirical evidence that choice sets are limited by bounded rationality even in an online setting where search costs are minimized.

The remainder of this paper is organized as follows. Sections 2 and 3 describe the setting and data used in our analysis, respectively. Section 4 presents our empirical approach to identify saliency effects in the data. Sections 5 & 6 present our results and Section 7 offers a few concluding remarks.

2 Description of Online Market

We use data from an online luxury fashion retailer, Net-a-Porter, selling top fashion brands such as Burberry, Dolce & Gabbana, Gucci and Dior. Founded in 2000, Net-a-Porter sells fashion, shoes and accessories to 170 countries.[9] The company sells almost exclusively to women, the majority of whom have a graduate degree, and claims that its average consumer has an (annual) household income of \$170,000 and expenditure on fashion is \$13,000. It claims 6 million unique users worldwide every month, with a third in the U.S. and 40% in the UK and the rest of Europe, with an average value of an order at \$500. The website is highly successful, with a bounce rate of 34.8% in 2015 and an average of slightly over 6 page views per visitor.[10]

Net-a-Porter is widely considered to have revolutionised retailing luxury fashion because from a customer's perspective it does away with the experience of shopping in an exclusive store and from the fashion label's point of view it dispensed with the need for expensive retail stores. To achieve this, Net-a-Porter undertakes efforts to raise confidence and reduce the risk in only luxury shopping by offering extensive product views (including videos, measurements of products and detailed product description), careful distilling of trends, and an efficient global courier delivery system, with 24 hour delivery service in London and New York.

There are three points about this retailer that make it a useful setting for examining consumer choice. The first is that the website provides recommendations for every product which are non-personalised. This is ideal for our analysis as we are able to avoid dealing with a large

[9]It was founded in 2000 as a small start-up and is now part of Richemont, a Swiss-based luxury conglomerate, which bought a 93% stake in 2010. Net-a-Porter had sales of Eur 550m last year and is now worth Eur 2.5 billion, roughly six times its value when Richemont invested in it.

[10]The bounce rate is the percentage of users that arrive on the website and leave without viewing a second page on the site. Net-a-Porter's bounce rate is comparable to that of other highly successful online luxury retailers such as Neiman Marcus or Mytheresa. Data from www.alexa.com.

amount of consumer heterogeneity present in most personalised recommendation systems. Second, the only other information provided on each product page, in addition to recommendations, are product attributes (image, price, description, dimensions etc). Importantly, the information does not include any signal on the underlying popularity of the product through reviews or sales-rank or any such instruments. This ensures that users are not choosing products based on their popularity on the web-site as such information is absent. Finally, the concept of the web-site is innovative, marrying both content and commerce, enabling it to attract a large volume of customers. In brief, the site allows us to examine choice across more than 15,000 products and 530 brands, in a setting where consumers are largely fully informed about product and brand attributes including prices and product recommendations are not tailored towards individual customers.

3 Data

The data were obtained from the Net-a-Porter website between May and August 2014. The main dependent variable that we use consists of information on additions to the shopping bag and wishlists by anonymous buyers which provides information on potential sales of products.[11] The data consists of several components:

Products: We parsed the entire set of available products from Net-a-Porter's product catalogue. The catalogue distinguishes between broad categories: clothing, bags, shoes, accessories, lingerie, sport, and beauty. There are a number of subcategories within the broader categories (for example dresses, pants, skirts etc. under clothing). The catalogue presents the products with a number of photos and basic product attributes including the price. Once a customer clicks on the product, a detailed description appears plus more photos and videos.

Product-level transactions: Net-a-Porter's online platform includes a feature called Net-a-Porter "Live" which provides real time data on product sales. The live data feed, updated every second, allows customers to see how many people around the world (and indeed in their particular location) are browsing the site with them, and what they are adding to their shopping bags and wishlists. The transaction-level data used in our analysis comes from this live ticker. That is, we have product-level information on all items that customers have added to their wishlists and shopping bags, which includes basic information on product attributes including brands and prices as well as the precise time when a customer made these transactions and her physical location. In an informal, confidential discussion with the

[11]U.S. and U.K. regulations identify several research categories that are considered exempt from Institutional Review Board oversight (see Office for Human Research Protections (January 15, 2009). "Code of Federal Regulations". hhs.gov. p. US 45 CFR 46.101. Also see http://www.research-integrity.admin.cam.ac.uk/research-ethics/guidance). Research involving the analysis of existing data and other materials if they are already publicly available, or where the data can be collected such that individual subjects cannot be identified in any way is not subject to such oversight. In our case, we have no information on individual buyers and the data used in our analysis are publicly visible information on the Net-a-Porter website. Hence, our analysis and all reported results do not reveal any information on individual users. Neither do we reveal any specific statistics on prices, characteristics, or shopping bag/wishlist additions of individual products or aggregate data on total shopping bag/wishlist additions across product groups or the entire website.

representatives from the company, we obtained information on the manner in which these data are presented and the implications for our analysis. First, while the ticker tape does not provide data on every potential transaction, it is a random sample of transactions, which is updated every 8-10 minutes. The main idea here is that viewers of the web page stay on the page for an average of 3 minutes or less and thus updating the list every 8-10 minutes allows enough variation. Second, it is possible that some additions to the shopping bag do not result in actual transactions; however, purchases cannot be completed without adding to the shopping bag. In brief, we have a random sample of potential purchases in these data. We interpret additions to a customer's shopping bag as a serious intention to buy the corresponding product (hence we refer to it as sales), whereas additions to the wishlist are interpreted as an intention to buy. While these transaction-level data are available per minute, we aggregate the data to daily intervals.[12] Table 1 summarizes the data from the live ticker. The table distinguishes between existing and new products (see next bullet point) as well as among products that are recommended by a new product and those that are not.

New products: Net-a-Porter launches new products three times a week: on Monday, Wednesday, and Friday. We identify all new products from the "What's New" category on Net-a-Porter's website, which lists all new arrivals. Supply factors largely determine the timing of these launches. Products are launched on the website as and when they are released by the product's producer to Net-a-Porter's warehouses.

Product recommendations *"you may also like"* (substitutes): In addition to detailed information on a given product, the customer is also provided with product recommendations under the **you may also like** header – see Figure 1. The products under the **you may also like** header form the recommendation sets used in our analysis. These are products that are very similar to the target product, they usually belong to the same product category (in our sample 99% of recommended products are in the same category as their recommending new products – see Table 2), a similar price range (the average price difference between recommended and new recommending products in our sample is close to zero – see Table 2), but not necessarily the same brand/designer (see Table 2). The number of recommended products differs by product. Unlike standard models of product referrals, such as Amazon, Net-a-Porter's recommended sets are not personalised (as in "if *you* like this, *you* will also like") and are simply potential substitutes (see Section 5.1 for more details).[13]

[12]There are three main reasons for aggregating at the daily-level even though the aggregation entails a loss of information. First, the major source of variation for our variable of interest, saliency shock, is at the daily level. As such the additional information contained at a lower level of aggregation is not particulary useful for our purposes. Secondly, minute level transaction data are highly volatile and our aggregation scheme enables us to reduce the impact of non-relevant microstructure effects that induce noise. Finally, daily aggregation reduces the amount of data used for analysis, allowing us to compute our estimates more efficiently. To put further structure on the level of our aggregation, in section 4.1.1, we specify the distribution of the daily transaction volume as a poisson process. Our dependent variable therefore, counts the number of times a shopping event has occurred during a daily time interval.

[13]The sets are not personalised because the aim of the company is to create a shopping experience akin to browsing a fashion magazine or shopping in a physical store. More recently, Net-a-Porter has created apps that allow targeting specific audiences through interaction on social media networks. However, our data pre-date this.

In conversations with Net-a-Porter we learned that substitutes are chosen through a combination of two tools: one that selects visually similar products[14] and another one that selects products with similar observable attributes. However, no attempt is made to customize the recommendation sets based on some (subjective) perception of product popularity. We also confirmed in our conversations with Net-a-Porter that recommendations are not chosen based on past or expected sales of either the recommending or recommended products. According to them, the goal of providing recommendations is mainly to suggest similar products (substitutes), similar to what a customer would experience in a brick and mortar fashion boutique.

Regional variation in recommendation sets: It turns out that there is some small variation in recommended sets across regions which we exploit in our empirical analysis (see Section 4.1.2). That is, in a few cases, the same new product recommends slightly different sets of products in different regions, say the U.S. and Europe. According to Net-a-Porter this is the result of the attribute matching tool placing different weights on a product's attributes in attempts to accommodate taste differences across regions. These product- and region-specific taste differences, termed as '*style*' by the merchandise team,[15] concern variation in preferences for designer labels and fashion type (classic vs contemporary etc.) and are fixed over a product's life-cycle. This means that any observed differences between recommendation sets across regions are due to region-specific market characteristics rather than product-specific demand trends.

Product recommendations *"how to wear it"* (complements): The website offers also product recommendations under a **how to wear it** header. These recommendations are products that can be worn in combination with the target product. Hence, usually these are products from other product categories (for example if the target product is a dress, **how to wear it** might show shoes, a bag, and earrings). As such we consider these products as complements as opposed to substitutes under the **you may also like** header.

International dimension: Net-a-Porter splits its offer into three geographical areas: the Americas (which includes the U.S.), International (which includes Europe), and Asia and the Pacific (which includes India and China). Part of our analysis relies on variation in the sets of recommended products across these areas (see Figure 2). To obtain the data for the Americas and Europe, we parsed the data from locations in the U.S. and the UK. Since the live ticker provides us with the location of customers, we can determine which set of recommended products a given customer in a given region was able to see.

[14]The Guardian, in a recent article labels the technology used here as "world-class image recognition technology", see http://www.theguardian.com/media-network/2015/may/14/net-a-porter-fashion-digital-revolution

[15]The company distinguishes between fixed effects of location, named *style* and time-varying *trends*, which are thought to be the same across locations. We were told explicitly that recommendation sets are not designed as per *trend*.

4 Reduced Form Empirical Specification and Identification

To establish the relationship between product saliency, through recommendations, and demand we begin by laying out a reduced form specification linking product sales to their saliency. Later in section 6 we provide more structure to the reduced form effects, by testing whether recommendations impact sales by expanding consumer consideration sets, that ultimately affect their demand for the product. Before doing that however, we use the reduced form specification to describe and motivate our identification strategy for consistently estimating the causal effect of saliency. The effect of product saliency on its aggregate sales can be written as:

$$y_{jt} = \alpha + \psi s_{jt} + \epsilon_{jt} \tag{4.1}$$

The dependent variable, y_{jt} is the total number of shopping bag or wishlist additions for product j during calender day t. Our main variable of interest is product saliency, s_{jt}, This variable is defined as follows: We denote s_{jkt} as an indicator for whether product j is included in the set of recommended products for a product k. We then sum this variable over the set of all products in the catalogue at date t giving us, $s_{jt} = \sum_{k,k \neq j} s_{jkt}$. The variable s_{jt} measures the intensity of saliency for an existing product j. The coefficient ψ measures the total impact of saliency on product demand and captures both the consideration and choice probability.

Least squares estimation of this equation will however produce a biased estimate of parameter ψ if products appearing in recommendation sets are endogenously selected. This occurs, for example, if products that have experienced high demand are targeted specifically through recommendations, or when similar cheaper products are systematically recommended as substitutes. Take for example the case of a best-selling handbag . An endogeneity problem would emerge if the retailer wants to draw attention to its popularity by recommending it frequently with other products in the same-category. In this case, the saliency parameter, ψ, would be upward-biased as a result of the positive correlation between unobserved popularity and the frequency of recommendation. The next section explains how we address the potential endogeneity issue.

4.1 Sources of Identification

4.1.1 New Product Arrival Shocks

To consistently estimate the effect of saliency we use identifying variation from the arrival of new products and their impact on existing products. Our identification strategy exploits two features of new products. First, new products are more salient on behalf of their novelty. Dedicated web-links and email based advertising to announce the arrival of these products increases their popularity at the time of their arrival. Second, most new products recommend other products that have already existed on the website but have not been recommended in these sets before. The in-stock products have a demand history that allows us to control for their latent popularity, thereby eliminating potential selection bias based on past sales.

We therefore make use of the spillover effects of the increased popularity of new products on the set of existing products that get recommended alongside them. We treat the arrival of new products as shocks to to existing products' saliency and use this to identify the saliency effect. In order to do this, we define the set of new products at every launch date as S. We denote $\widehat{s}_{jnt}^{N}$ as an indicator for whether product j is included in the set of recommended products for a **new product** n (where $n \in S$) launched at date t. We then sum this variable over the set of all possible new products launched at date t giving us, $\widehat{s}_{jt}^{N} = \sum_{n \in S, j \neq j} s_{jnt}$. The treatment variable $\widehat{s}_{jt}^{N}$ measures the intensity of saliency for an existing product j due to the arrival shock of new products at time t.

$$y_{jt} = \alpha + \sum_{\lambda=-\tau}^{\Gamma} \psi_{\tau-\lambda} \widehat{s}_{j(\tau-\lambda)}^{N} + \mu_j + \gamma_t + \epsilon_{jt} \tag{4.2}$$

We use a finite distributed lag model to estimate our model allowing saliency effects on product demand to last up to Γ days. For each date, we measure the length of the non-overlapping effect window λ by the number of days preceding the arrival shock to the number of days following the arrival shock. The reason for including leads on the treatment variables is to mitigate concerns about potential selection bias due to anticipation effects, if top (or low) selling existing products are endogenously chosen to be part of recommendation sets. We include product fixed effects μ_j to absorb product specific heterogeneity. We also accommodate different time trends in product demand by incorporating calender day fixed effects.

To estimate Equation 4.2, we use a fixed effect poisson model, as our dependent variable – the daily *count* of total shopping bag or wish-list additions – follows a poisson distribution.[16] Denote the vector of all explanatory variables by $\mathbf{x_{jt}}$ and the vector of all coefficients as $\mathbf{\Phi}$. In a fixed effect specification, the conditional likelihood is conditioned on the sum of the outcomes over the product-specific panel dimension (T_j) (Cameron and Trivedi 2013). For inference, we use cluster-robust standard errors, where each cluster is a product, to account for product specific serial correlation in ϵ_{jt} using the formula derived in Wooldridge (1999).

4.1.2 Regional Variation in Recommendation Sets

To further strengthen our identification strategy, in addition to employing product arrival shocks, we exploit regional variation in recommendation sets. One concern with our previously described specification is that new products and products that they recommend share

[16] As noted before, we specify the distribution of the daily transaction volume as a poisson process deriving from the underlying second-level transaction generating data. Our dependent variable counts the number of times a product has been added to the shopping basket or wishlist during a daily time interval. We assume implicitly that this is a process with independent increments – this is a justifiable assumption given that the website does not display product-specific shopping trends (cumulative or daily) and that the customers shopping on this website are not explicitly related in any way and are unable to fully observe each other's purchases.

some correlated attributes. If the attributes that make a new product sell well are correlated with the attributes that determine which existing product are recommended, then those existing products will also sell well post the introduction of the new product, not because of the recommendation but because of the shared attributes. Even though we control for the past sales record of each recommended product in the previous specification, we further strengthen our identification to resolve the issue of correlated effects by exploiting regional variations in the composition of recommendation sets. This allows us to include a new product fixed effect and absorb all sales-enhancing correlated effects. In addition the strategy allows us to difference out product-by-time unobservable characteristics.

As explained in Section 3, Net-a-Porter provides different recommendation sets across regions for a few products . These differences are a result of regional taste differences that are specific to a product but do not vary over its life-cycle. With that in mind, we decompose the region-specific residual (for each region R), ϵ_{jt}^{R} into the following components:

$$\epsilon_{jt}^{R} = \omega_{jt} + \mu_{j}^{R} + \upsilon_{jt}^{R} \tag{4.3}$$

where ω_{jt} denotes the time-varying product specific unobservables that are common across all regions, μ_{j}^{R} is the time-invariant product specific unobservable that differs across regions and υ_{jt}^{R} is the time-varying product specific unobservable that differs across regions. For example, in our set-up, μ_{j}^{R} captures fixed regional differences in preferences for each product while υ_{jt}^{R} captures the differential shift in these preferences.

We now describe how, the regional variation in recommendation sets for the same new product allows us to difference out the relevant components of the overall residual. To begin with, we normalize our time variable to event-time days and restrict our analysis to -3/+3 days of product j receiving the saliency shock (day 0). We consider the two regions in which most of the transactions occur - America and Europe - and introduce an additional subscript n which indexes the overall product recommendation set associated with a new product. Therefore y_{jnt}^{R} is the total shopping bag additions for product j recommended by new product n at time t in region R. For each region, we define the treatment variable, $\widehat{T}_{jn}^{R}$, as an indicator taking the value one if product j was recommended by new product n in region R. $Post_t$ is an indicator for the time period following the saliency shock i.e. day 0 to day 3.

$$\begin{aligned} y_{jnt}^{AMR} &= \beta_1 \widehat{T}_{jn}^{AMR} + \beta_2 (T_{jn}^{AMR} \times Post_t) + \mu_j + \gamma_t + \underbrace{\omega_{jnt} + \mu_{in}^{AMR} + \upsilon_{jnt}^{AMR}} && (4.4) \\ y_{jnt}^{EUR} &= \beta_1 T_{jn}^{EUR} + \beta_2 (T_{jn}^{EUR} \times Post_t) + \mu_j + \gamma_t + \underbrace{\omega_{jnt} + \mu_{jn}^{EUR} + \upsilon_{jnt}^{EUR}} && (4.5) \end{aligned}$$

Now, we can net out the time-varying product specific unobservables that are potentially correlated with a product receiving a saliency shock by taking a difference of Equations (4.4-4.5):

$$y^{\star}_{jnt} = \beta_1 T^{\star}_{jn} + \beta_2(T^{\star}_{jn} \times Post_t) + \underbrace{\mu^{\star}_{jn} + \upsilon^{\star}_{jnt}} \tag{4.6}$$

where $y^{\star}_{jnt}$ denotes the *difference* in demand between America and Europe for product j recommended by new product n at time t. Similarly, $T^{\star}_{jn}$ denotes the difference in treatment status of product j, i.e. whether it is recommended by new product n, between America and Europe. Note that our differencing strategy is, implicitly, only relevant for products that were exclusively recommended in either of the two regions; as a result we discard all products that were recommended both in America and Europe. For ease of interpretation, we also recode the variable, $T^{\star}$ to take the value 0 if the product was recommended in Europe but not in America (instead of -1 as the differencing suggests).

It is easy to see that first-differencing equation 4.6 allows us to absorb the time-invariant regional unobservables for each product j, $\mu^{\star}_{jn}$, that are crucial in determining the allocation of products to different sets across regions . Further, to deal with the issue of correlated attributes between the new and recommending product we rely on the within (new) product regional variation in recommendation sets. We include, in the differenced equation, a fixed effect ($\mathbf{B_n}$) for each new product n that recommends different products in different regions:

$$\Delta y^{\star}_{jnt} = \beta_2 \Delta(T^{\star}_{jn} \times Post_t) + \mathbf{B_n} + \Delta\upsilon^{\star}_{jnt} \tag{4.7}$$

Including the fixed effect acts as a synthetic control (Abadie et al. 2010), allowing us to compare the demand differential between product j and product k that are recommended by the *same new product n* but in different regions. Our identifying assumption for the difference-in-difference specification is that conditional on being recommended for the same product, product j and product k experience similar trends in product sales before being recommended by the new product.[17] In section 5.6 we empirically test this common trends assumption and show that it is validated in our data.

5 Results: Impact of Saliency on Aggregate Sales

5.1 Identifying Conditions

Our identification strategy rests on two assumptions that we can test empirically. First, our strategy requires that new products produce saliency shocks for recommended products. To show that new products are themselves highly salient, Figure 4 plots the novelty effect for new products. The figure shows that new products are highly popular upon arrival and this effect declines over the week. The largest effects are observed over the first four days with the effect tapering off by the sixth day following the arrival. The figure suggests that new

[17] For instance, a concern could be that products chosen for recommendation in the U.S. started to experience a higher sales trend, before being recommended, compared to a non-recommended product in Europe.

products attract an enormous amount of consumer attention and demand immediately when they are launched on the platform.

Second, our empirical approach relies on the fact that recommendation sets for new products are not endogenously selected (although we relax this assumption when we exploit regional differences in recommendation sets). To investigate this assumption in the data with regard to observable product characteristics, Table 3 shows average prices for recommended (by new products) and non-recommended products across the different product categories. On average, we do not find any significant differences at reasonable levels of statistical significance. We also test whether products that are recommended by new products are subject to different demand prior to being recommended. Figure 3 shows the empirical distribution of the difference in shopping bag additions during a 3, 5, and 7 day time period prior to being recommended by a new product and shopping bag additions during the same time period of all other products.[18] We compute these differences comparing all products (upper graph) and comparing only products within product categories (lower graph). The red colors show that in the overwhelming majority of cases there are no statistical differences in demand for products that are recommended subsequently (3, 5, or 7 days later) by a new product and all other products. Note that while our basic results rely on this conditional independence assumption, our estimations that use differences in recommendation sets across regions difference out product-by-time unobservables.

5.2 Baseline Specification

For all our specifications and results, we refer to "saliency" for product j as the number of new products recommending this product at any given point in time. Table 5 reports estimates for the average effect of saliency on total shopping bag additions on the day it received the shock. First we discuss estimates for the effect of a product being new. Column (1) shows that, upon arrival, new products have on average 96% more shopping bag additions compared to existing products. This confirms our assumption that new products are highly popular and are likely to generate spillover effects from their popularity. We now examine the effect of being recommended by a new product (see Table 4 for descriptive statistics). Column (1) shows that a one unit increase in saliency, i.e. an additional new product recommendation, increases the total number of shopping bag additions by approximately 6%. To incorporate potential anticipation effects we include the forward lag of saliency in Column (2) and find that our results are still robust to this addition. Further, the coefficient on the forward lag is negative but statistically insignificant suggesting that products exposed to a saliency shock did not experience a differential demand trend prior to receiving the saliency shock. In Column (3) we split the saliency effect between existing products and new products. As described in the data section, new products recommended a mix of existing and (other) new products. We find that the effect of saliency is large and significant for existing products which see a 5.5%

[18] Each bar displayed in Figure 3 corresponds to the arrival date of a new product, since only the arrival of a new product leads to the recommendation of an existing product.

increase in their sales on the day that they are recommended by a new product. However, this effect is close to zero for new products receiving the saliency shock implying that the novelty effect dominates the sales of new products upon arrival and that there are no added affects of recommendations. Finally in Column (4) we control for the lagged effect (up to 2 weeks) of being a new product addressing the concern that saliency shocks might be picking up lagged new product effects if it were the case that lagged new products were likely to receive the saliency shock. We find that our result is robust to including this control and that the saliency effect is independent of lagged new product effects.

Until now we have focused on the immediate short term effect of a saliency shock. To assess whether these saliency effects persist over the days following the arrival of the product, we report results from estimating the finite distributed lag model presented in Equation 4.2. Figure 5 plots the disaggregated saliency effects for each day following the shock along-with their confidence intervals. The figure shows a large increase in total purchases for salient products on the day they receive the saliency shock (6% increase) with the effect positive but declining over the subsequent few days. On average an additional unit of saliency results in a 3-5% increase in total purchases over the three days following the shock. Such a pattern of effects, with a prominent spike in a product's sales on the day that they are recommended, is consistent with the "attention" based explanation that products receiving a saliency shock could have been previously overlooked by consumers. The lack of persistent long-term effects is incompatible with a selection-based explanation which posits that existing products are endogenously selected to be part of new product recommendation sets in anticipation of their (higher) future sales.

In figure 6 we break down the daily effects of saliency for existing and new products. The results are mixed. Figure 6(a) shows that existing products see a large increase in their sales on the day that they are recommended by a new product but this effect disappears on day 2, subsequently picking up again over days 3 and 4. In contrast, we find no effect of saliency for new products on the day they are launched *and* recommended by other new products (figure 6(b)) but we find positive and significant effects following the day of the shock. Our results indicate that while the novelty feature of a new product clearly dominates its sales on the day of its launch, the saliency effect starts to play a role in increasing its sales once the novelty effect starts wearing off, over the subsequent few days. This means that new products that were recommended by other new products are able to maintain a competitive edge in the days following their launch compared to new products that are not recommended by other new products.

Since our data also contain information on consumers' "wishlists," we undertake the same analysis with total wishlist additions (per day) as a a dependent variable. Table 6 reports these results. Although additions to wishlists are more noisy, we find strikingly similar results for the effect of saliency. Across all specifications we find that a one unit increase in saliency, i.e., an additional new product recommendation, increases the total number of wishlist additions by approximately 6%. This result is robust to controlling for lagged new product effects and anticipation effects. Further we find strong novelty effects with new products experiencing almost a 118% increase in wishlist additions compared to existing products. We find close to zero effects of saliency shocks for new products on the day they receive the shock (Columns (3) and (4)).

Finally, Figure 13 plots coefficients when we use recommendations of complements (*"how to wear it"*) instead of substitutes (*"you may also like"*). The results are similar to those obtained for substitutes in Figure 6; focusing on the effect on existing products, we see a significant, positive effect of the saliency shock on demand on the day of the saliency shock, with the effect pattering out within the first three days. The fact that the coefficient is smaller in magnitude than in the case of substitutes is what we would expect if most consumers choose between substitutes rather than switching to or adding complements to their shopping bags. The effect of the saliency shock on other new products shown in the lower plot is much more lasting, we continue to observe a positive impact up to seven days following the shock.

5.3 Robustness

To assess the robustness of our results to lagged novelty effects, we include the two-week lag of whether a product was new in the baseline specification. Figures 7(a)–7(b) show that the results, both for existing and new products, are robust to the inclusion of this control. We also assess whether our results are sensitive to introducing differential anticipation effects between products that received a saliency shock in the previous week compared to products that did not. *A priori*, one might expect that products that received a saliency shock in the previous week have an upward trending sales curve that makes them more likely to receive another saliency shock. If this were the case, then we would be picking up lagged saliency effects of high-selling products, confounding our estimates of current saliency shocks.

Figure 8(a) plots the results with the anticipation effects split between products that received a prior saliency shock and those that did not. We see clearly that there is almost no difference in the anticipation effects between these two types of products and that the overall effects are close to zero. Products that received a saliency shock had no differential demand trend 3 days prior to the event. We conduct the same analysis for total wishlist additions as a dependent variable. Similar to the results for shopping bag additions, we find that the saliency effects for consumers' wishlist additions are robust to controlling for lagged new product effects (Figure 11) and differential anticipation effects (Figure 12). We conclude therefore, that our results are not affected by including the various controls described above.

5.4 Attention and Price Effects

We now examine heterogeneity in saliency effects across products. The first dimension of heterogeneity we explore is in the size of recommendation sets. For each recommended product j, we compute the display size of the set in which product i was recommended; it equals the total number of other products that were also recommended alongside product j by a new product n. The overall size of the recommendation set may matter if consumers have limited attention and can only focus on a restricted number of products at a time. As a result, products that are recommended in smaller sets may receive more attention, increasing their sales, compared to products recommended in larger sets. To test the display size effect, we include an interaction of the saliency shock to product j and size of the set in which it was recommended. Figure 9 shows results from this specification. We find substantial, negative and significant, display size effects. An additional product in the recommendation set reduces total purchases by 4%. To gauge the magnitude of this effect, we note that the average size of the recommendation set in our sample is 7.5. On average, therefore, products recommended with 7 other products see an increase in their sales by about 6%, similar to the results found in our baseline specification. The maximum effect of saliency is experienced by products recommended in sets of 1-3.

We now turn to exploring the price sensitivity of salient products. To examine this, we interact the saliency shock with the difference in price between product j and the new product n which recommends it. Our null hypothesis is that consumers are less likely to respond to price differences if product recommendations serve only to improve the saliency of a product, thereby drawing consumers' attention. Figure 10 plots the results of the interactions. We fail to reject the null for the interaction effects both on the day of the shock and subsequently. The coefficient on the interaction terms is close to zero suggesting that consumers ignore variation in price differences across recommended products and are influenced only by the number of other competing products.

5.5 Spillover Effects of Saliency

So far, we have measured the direct effect of a saliency shock on products that are recommended by newly launched products. To the extent that these products also recommend other products, there could exist substantial spillover effects of the saliency shock that potentially bias our estimates downwards. To explore the presence of spillover effects, we build a network of recommendations that allow us to vertically trace the impact of the saliency shock, originating from newly launched products.

We measure, on a given day, the path distance between a product and a new product in the recommendation network. For example, products that were directly recommended by a new product have a one degree separation and are identified by the dummy variable, D^1. Further, products that are recommended by degree 1 products have a two degree separation between themselves and the new product and are identified by the dummy variable, D^2. In a similar way we identify degree 3 products (D^3). Note that the variables that identify the degree of a product are mutually exclusive, in the sense that products are identified by their closest degree of separation even if they can be recommended recursively through the network.

$$
\begin{aligned}
y_{jt} \quad = \quad & \alpha + \sum_{\lambda=-\tau}^{\Gamma} \psi_{\tau-\lambda}(\widehat{s}^{N}_{j(\tau-\lambda)} \times D^{1}_{j(\tau-\lambda)}) + \sum_{\lambda=-\tau}^{\Gamma} \psi_{\tau-\lambda}(\widehat{s}^{N}_{j(\tau-\lambda)} \times D^{2}_{j(\tau-\lambda)}) \\
& + \sum_{\lambda=-\tau}^{\Gamma} \psi_{\tau-\lambda}(\widehat{s}^{N}_{j\tau-\lambda} \times D^{3}_{j(\tau-\lambda)}) + \mu_j + \gamma_t + \epsilon_{jt} \qquad (5.1)
\end{aligned}
$$

In this specification, $\widehat{s}^{N}_{j(\tau-\lambda)} \times D^{1}_{j(\tau-\lambda)}$ measures the total number of new products recommending product j, i.e., it represents the intensity of saliency for products that are directly recommended by new products (degree 1). The indirect spillover effects are captured by the variables $\widehat{s}^{N}_{j(\tau-\lambda)} \times D^{2}_{j(\tau-\lambda)}$ and $\widehat{s}^{N}_{j(\tau-\lambda)} \times D^{3}_{j(\tau-\lambda)}$. $\widehat{s}^{N}_{j(\tau-\lambda)} \times D^{2}_{j(\tau-\lambda)}$ measures the total number of degree 1 products (those directly affected by the saliency shock) recommending product j for all products at a two degree separation from any new product; $\widehat{s}^{N}_{j(\tau-\lambda)} \times D^{2}_{j(\tau-\lambda)}$ measures the total number of degree 2 products (those indirectly affected by the saliency shock) recommending product j for all products at a three degree separation from any new product.

Table 7 reports results from including spillover effects. The first row of the table presents our baseline results, where we do not account for spillover effects. The subsequent rows report results on both direct and indirect effects of the saliency shock. We find that our baseline results are largely unchanged by the inclusion of the spillover variables. As expected, there is a slight increase in the magnitude of the effect, from 7.4% to 8.2%, after accounting for spillover effects. A comparison of the direct and indirect effects of the saliency shocks reveals that the effects of the saliency shock are strongest for products recommended directly (at degree 1 separation) by new products. However we also find significant (positive) spillover effects. Products recommended by saliency-shock-affected products also see an increase in their sales on the day that new products are launched, but to a much lesser extent (a 2% increase). The spillover effects are limited to products at a degree 2 separation from new products. We find no significant effects for products that are at a three degree separation from new products.

5.6 Exploiting Regional Variation in Recommendation Sets

We have presented a range of results starting with the baseline specification and extensions in various directions. Table 8 shows the coefficients of interest for all the different models. The table highlights how consistent our results are across specifications – we see a large positive coefficient between 0.74 and 0.82 on the day an existing product receives a saliency shock, with this positive effect lasting for three days. The table also shows that these findings are unaffected by accounting for the price difference between recommending and recommended products. At the same time, we find that the larger the recommendation set, the smaller the saliency effect. These results paint a consistent picture, suggesting that products that have been available to consumers on the platform experience a large surge in demand when they are made more salient through a recommendation by highly salient products.

A concern with all the results presented in Table 8 is that the probability of receiving a saliency shock has an underlying correlation with future sales. In all our results we have shown that, on average, products receiving a saliency shock did not experience a differential demand trend from non-saliency shock products, prior to receiving the shock. In this section, we employ a tight and robust specification that differences out product-by-time unobservables and allows us to consistently estimate the saliency effect. As explained in Section 4.1.2, we exploit regional variation in the composition of recommendation sets to identify the saliency effect.

To begin with, we conduct a simple placebo test that illustrates our identification strategy. Figure 14 plots the results for saliency effects from our baseline specification (for existing and new products) and additionally reports results from a counterfactual exercise built around regional differences. We construct the counterfactual in the following way: we identify two sets of products – those receiving a saliency shock in America and those receiving a saliency shock in Europe. Next, we examine whether products that received a saliency shock exclusively in Europe, i.e., they were recommended by a new product (which itself was launched globally) only in Europe, increased in any way their American sales. Since consumers in America are not able to view these products as salient in their region, we should expect no change in the products' American sales if there was a pure saliency/attention effect driving up sales. Instead, if there were underlying time varying product trends for salient products, for example if salient products coincided with fashion trends picked up by their similarity to new products, then our hypothesis would be rejected and our identification would stand compromised. Figure 14(c) shows that this is not the case and that products made salient exclusively in Europe saw no change in their American demand. All effects, short-term and long-term, are close to zero.

Having described the essence of our identification approach, we now proceed to obtaining consistent effects for saliency using this strategy. The objective of our exercise is to estimate treatment effects, described as whether a product is made salient in America, on the sales differential between America and Europe. In estimating equation 4.7, we obtain estimates that – conditional on fixed effects for each new product launched globally – are independent of unobserved i) time-varying product differences, ii) time-invariant regional differences for each product and iii) time-varying regional differences for each product. Before reporting the results, we test the pre-event, unconditional difference in the America-Europe sales differential and prices of both treatment and control products. Table 9 shows that there are no statistically significant differences in the price and sales of treatment and control products before the saliency shock hits. These results suggest that even without conditioning on new product fixed effects, there are hardly any differences in the characteristics and outcomes of treated and untreated products. In contrast we find a large and significant difference in the post-saliency-shock sales of treatment and control products suggesting a positive saliency effect. In addition we test the common-trends assumption, implicit in our double difference-in-difference strategy. Figure 15 plots the (predicted)[19] difference in sales between America and Europe on the y-axis and event time on the x-axis. In both the figure and following results table, we undertake a *within new-product* comparison. This means we compare products that receive a saliency shock in America (treatment) by a new product n with a similar product (control) that is also recommended by n but only in Europe and not in the Americas. While we estimate equation 4.7 over a daily time interval of {-3. +3} days, the figure is extended to 12 hour (half-day) intervals over the same sample range.[20] The figure shows that both control (plotted in gray) and treatment (plotted in black) products have a declining sales curve but treatment products lie slightly below control products; however this difference is not statistically significant. On the day that treatment products receive their shock in America, their sales differential increases by a magnitude of almost two in favor of America. Following the event, the product continues its declining trend but the boost in its sales on the event day puts its sales curve on a higher level compared to control products reversing the pre-event trend gap.

Table 10 now reports results from our double difference-in-difference strategy. All columns condition on new product fixed effects. Column (1) estimates a simple difference-in-difference equation and retains the base treatment effect to show that the baseline difference between treatment and control products is statistically insignificant. In Column (2) we estimate the double difference-in-difference equation 4.7. We find that products recommended by a new product in America see a 13% increase in their American-Europe sales differential over the 4 event days, compared to similar products recommended exclusively in Europe by the same new product. The magnitude of the effect is larger than the effect obtained in our baseline specification. In column (3) we examine whether the treatment effects differ by size of the recommendation sets. We find a 2.5% decrease in sales differential with the inclusion of an additional product in the recommendation set. This effect is slightly smaller compared to what we obtained in our baseline specification but given the differences in the average size of the recommendation sets between our different samples, we conclude that the effect is largely similar. Finally, column (4) breaks down the effect by event day. As expected, we find a large, positive effect on the day of the event (10%) and surprisingly large effects sustained over the days following the event. In this sample, the effect of the saliency shock is largest (17% increase in sales differential) on the second day following the event.

6 Mechanisms and Structural Effects of Saliency on Consumer Demand

The reduced-form analysis has shown that a product's saliency has a positive and significant effect on its sales. Yet, our results from this analysis cannot distinguish whether this effect

[19] We use predicted difference in sales by regressing actual sales differentials on a fixed effect for each new product recommendation set. This allows us to test the common trends assumption conditional on $\mathbf{B_n}$ as required by our identification strategy.

[20] Note that since we require information on sales prior to the event, by construction, we are only able to examine saliency effects for existing products.

occurs through an increase in a consumers' consideration for this product or through consumers' implicit preference for saliency (thereby affecting their choice). Indeed, we have assumed that the effect of saliency on sales works through consumers' consideration sets in order to interpret our results as evidence for limited attention.

In this section, we verify this assumption and estimate the two-stage probabilistic choice multinomial logit model (PCMNL). This approach allows us to test explicitly how saliency affects a consumer's consideration and choice and hence whether consumers make choices under limited attention. We also want to distinguish any such limited attention effect from consumer preferences for saliency.

6.1 Probabilistic Choice Multinomial Logit Model

The difficulty in the literature is distinguishing whether a product is not consumed because it has no utility to the consumer or whether the consumer is simply unaware of it because of scarce attention. The latter would imply that consumers do not take into account all the alternatives available on the website in making their choices but reduce them to a smaller (manageable) set. A consumer might be able to see all the alternatives available to her but only evaluates a subset of them, the choice or consideration set, and makes her final choice by maximising her preferences over this set.

This process is described by Manski (1977) in his econometric formulation of choice behaviour that allows for sequential decisions with heterogenous choice sets. Manski (1977) proposes random utility models of choice, where choice sets are probabilistic in nature and final choices are conditional on this choice set. Within this class of model, we adopt the random constraint-based approach of Swait and Ben-Akiva (1987) where a product is excluded from the choice set if its consideration utility is lower than some threshold consideration utility level. As described by Başar and Bhat (2004), since this threshold utility level is not observed by the econometrician, the exclusion of a product from the choice set becomes probabilistic. Note that in our framework, we do not explicitly model the consumer search process. In general, consumers may face heterogenous search costs that are fixed across all products (De los Santos et al. 2012) or vary by product (Kim et al. 2010) but our consideration stage specification is agnostic about this. In this sense, our results for the consideration effect can be viewed as a mixture of increasing awareness and lowering search costs. [21].

We consider the probability that alternative j is considered by consumer i at any time t, where

[21] Honka et al. (2014) model all three stages of a consumer's decision process, awareness, consideration and choice and find that advertising serves to mainly increase awareness for a product. Their data identifies the list of options considered by the consumers during their search process.

t represents a calender day.[22] This probability can be written as:

$$C_{ijt} = \frac{1}{1 + e^{-(\phi' \mathbf{w_{ijt}} + \psi_1' s_{jt})}} \tag{6.1}$$

where $\mathbf{w_{ijt}}$ is a column vector of observed attributes for user i and alternative j at time t and ϕ is a corresponding column vector of coefficients which provide the impact of attributes on the consideration probability of alternative j. Our variable of interest, product saliency is captured by s_{jt} which is defined as the number of recommended sets by a new product that product j appears in at time t of the corresponding new products' launch. The coefficient ψ_1 measures the impact of saliency on the consideration probability of alternative j. An important identifying condition that we require for analysis is that, conditional on salience (and other included attributes, $\mathbf{w}$), the probability of consideration is independent across alternatives. While slightly restrictive, we justify this assumption in our data based on the fact that all products in each sub-category we analyse are fully substitutable. In addition, apart from saliency, there is very little menu dependence amongst alternatives i.e, each alternative is presented without any special distinguishing aspects. Allowing for dependence across alternatives in unfeasible in our context as it would yield no observable restrictions on the choice data with which to identify the consideration set as shown by Manzini and Mariotti (2014).[23]

The overall probability of a choice set c_t at time t for user i is given by:

$$P_{it}(c_t) = \frac{\prod_{j \in c_t} C_{ijt} \prod_{k \notin c_t} (1 - C_{ikt})}{1 - \prod_{j=1}^{J} (1 - C_{ijt})} \tag{6.2}$$

Note that the denominator is normalized to remove the "empty" choice set. It is also assumed that the randomly-distributed threshold for each alternative is independent of the threshold values of other alternatives. Conditional on the choice set, a consumer chooses product j at time t based on the following multinomial logit formulation, as:

$$P_{ijt}|c_t = \frac{e^{\boldsymbol{\beta}' \mathbf{x_{ijt}} + \psi_2 s_{jt}}}{\sum_{k \in c_t} e^{\boldsymbol{\beta}' \mathbf{x_{ikt}} + \psi_2 s_{kt}}} \quad \text{if} \quad j \in c_t \tag{6.3}$$

$$= 0 \quad \text{if} \quad j \notin c_t \tag{6.4}$$

where $\mathbf{x_{ijt}}$ is a column vector of exogenous variables that affect the probability of selecting a product conditional on a consumers choice set, $\boldsymbol{\beta}$ is a column vector of associated coefficients

[22] Implicitly we make the simplifying assumption that any user i considers purchasing only one unit of a product per day. This is not, however a restrictive assumption, and we can easily re-write the model in terms of a user considering to purchase a product at any given fraction of time.

[23] Manzini and Mariotti (2014) also show that menu dependence would lead to the preference relation being entirely unidentified. They cite evidence, from different contexts, that support no or weak menu effects.

and ψ_2 measures the impact of saliency on the choice probability of alternative j conditional on considering it. Given the conditional choice probability, the unconditional probability of choice of alternative j can be written as:

$$P_{ijt} = \sum_{c_t \in G} (P_{ijt}|c_t) \cdot P_{it}(c_t) \tag{6.5}$$

where G is the set of all non-empty subsets of the comprehensive choice set of all product alternatives, i.e., it includes each possible choice set, a total of $(2^I - 1)$ elements where I is the total number of products in the market. We estimate the consideration and choice stage parameters by iterating over all possible sets and maximizing the following log-likelihood function:

$$L(\phi, \psi_1, \beta, \psi_2) = \sum_i \sum_j y_{ijt} \cdot \log P_{ijt}(\phi, \psi_1, \beta, \psi_2) \tag{6.6}$$

where y_{ijt} is a dummy variable taking the value 1 if individual i chooses product j and 0 otherwise.

We can compute the disaggregate elasticity effects based on Başar and Bhat (2004). We define δ_{ijt} as a an indicator for whether the choice set c_t contains product j. Then, the probability, B_{ijt}, that the individual's choice set includes product alternative j is:

$$B_{ijt} = \sum_{c_t \in G} \delta_{ijt}^{c_t} P_{it}(c_t) = \frac{C_{ijt}}{1 - \prod_k (1 - C_{ikt})} \tag{6.7}$$

Now, consider the impact of attribute s_{it} that appears both at the consideration stage and choice stage. The overall self-elasticity (probability of choosing product j) and cross-elasticity (probability of choosing product k) of demand, with respect to the saliency of product j is given by (Başar and Bhat 2004):

$$\eta_{s_{jt}}^{P_{ijt}} = \left[\underbrace{(1 - B_{ijt})\psi_1}_{\text{Consideration}} + \underbrace{\frac{1}{P_{ijt}} \sum_{c_t \in G} \left\{ (P_{ijt}|c_t)(1 - P_{ijt}|c_t) P_{it}(c_t) \psi_2 \right\}}_{\text{Substitution}} \right] s_{jt} \tag{6.8}$$

$$\eta_{s_{jt}}^{P_{ikt}} = \left[\underbrace{\left\{ \frac{1}{P_{ik}} \sum_{c_t \in G} (P_{ikt}|c) P(c_t) \cdot \delta_{ijt}^{c_t} - B_{ijt} \right\} \psi_1}_{\text{Consideration}} + \underbrace{\frac{1}{P_{ikt}} \sum_{c_t \in G} \left\{ (-P_{ijt}|c_t)(P_{ikt}|c_t) P_{it}(c_t) \psi_2 \right\}}_{\text{Substitution}} \right) s_{jt} \tag{6.9}$$

The first term in both expressions is the consideration probability, which captures the impact of a change in the attribute on the consideration of product j while the second term represents the substitution probability at the choice stage conditional on product j being available in the choice set. The total effect of saliency depends, therefore, on the consideration probability for product j as well as its ultimate choice probability from amongst a set of considered alternatives.

In our setting, the probabilities of the choice set $P_{it}(c_t)$ are described by the random arrival of new products that in turn highlight (or make salient) a subset of older products, or additions to the (unobserved) choice set. These salience shocks to the existing full set of alternatives draw attention to the subset of existing alternatives highlighted by recommendations from the new arrivals. These shocks are not related to consumers' attributes, since they are not personalised. A particular advantage of our approach is that we do not need to know the exact choice set formulated by consumers but focus on the potential additions to the choice set created by the increased salience of a subset of products generated by new product arrivals.

Finally note, from the cross-elasticity expressions, that the PCMNL model does not display the Independence of Irrelevant Alternatives (IIA) feature of the multinomial/conditional logit model. The cross-elasticities in the PCMNL model depend on the probability information for *both* products j and k. This means that the cross-elasticities will be different across all alternatives.

6.2 Results: Effects of Saliency on Consumer Consideration and Choice

To estimate this model, we use transaction-level data on each consumer's purchase. Apart from consumer's geographic location, we do not identify any other consumer attributes. We thus treat the data as a pooled cross-section as our data does not track individual consumers over time. We are nevertheless able to identify the time period over which a given consumer visited the website and made a purchase decision. This allows us to include product level attributes in our models that vary over time and across consumers, such as saliency and novelty.

The reduced form analysis indicates that the effects of saliency last for a maximum of 4 days. Taking this into account, we define the saliency variable as a dummy variable taking the value one if the product was recommended by a new product at time t and over the 5 days following it. We define the variable "new product" in a similar way.[24]

Next in order to make our computation feasible, we estimate our model on a sub-set of data, i.e., instead of using information on the entire catalogue of products (that contains thousands of options), we narrow down on a sub-category of products, "Travel Bags", that contains only

[24]Note that although these new products are unavailable to consumers before they were launched, we still include them in our estimation. This is un-problematic because the the inclusion of the consideration stage, implicitly allows for some options to be irrelevant for some consumers (for example, those who visited the website when the product was not yet launched). For a more formal result, see Crawford et al. (2015) who discuss how the non-availability of products can be incorporated into a demand model with unobserved choice sets.

12 options.[25] For comparison, and to show that the results are not specific to the chosen product category, we also estimate the model for the sub-category "Watches" (with 25 options), but choosing only the 10 most purchased products that account for over 80% of sales. In addition to computational feasibility, we choose these two categories because products within these categories are highly substitutable.[26] For all specifications, we focus only on shopping bag additions and not on a consumer's wish-list as the latter provides a noisy measure of a consumer's actual choice.

Table 11 reports estimates from the the PCMNL model for both sub-categories (travel bags and watches). The majority of products in both sub-categories (more than 80%), receive a saliency shock at least once and vary only in the timing of receipt. Columns (1), (3) and (5) report coefficients from the consideration stage. We find that product saliency has a strong, positive and significant effect on consideration. Consumers, in our context, do not appear to be price-sensitive, but display a preference for choosing new products. While we find a small, positive and significant effect for saliency at the choice stage (Column (2)), this effect disappears when we disaggregate it for existing and new products. Column (3) shows that, for existing products, saliency has an insignificant effect on the probability of purchase in the choice stage. On the other hand, consumers are less likely to choose new products that are recommended (conditional on consideration), perhaps because they value new products that are more unique and less substitutable. Overall, taken together (consideration and choice), the aggregate marginal effects for saliency indicate an average increase in sales of 3% (travel bags)to 15%(watches) after being recommended by a new product (thereby increasing saliency). These effects are consistent with the results obtained from the reduced form analysis where we find that saliency increases product sales on average by approximately 6%. The reason why we find a lower marginal effect for travel-bags compared to watches is because most products in the travel bags category have a high consideration level in the market. Our model estimates imply that the average share of consumers who consider a product within the category travel-bags is a relatively high 78% compared to a low 49% for watches.[27]

Table 12 presents a sample of estimated own- and cross-price elasticities.[28]. Each entry j, k, where j indexes row and k column, gives the elasticity of product j with respect to a change in the saliency of k. Note that the products are labelled according to their sales rank (1

[25]For this reason, we are also unable to include product dummies in our specification. The inclusion of product dummies (in both the consideration and choice stage) makes the likelihood highly non-convex, resulting often in the non-convergence of our estimator. However we verify that the results from the MNL model are not sensitive to the inclusion or non-inclusion of product dummies. Counterfactual sales shares from either specification in the MNL are approximately the same.

[26]This is unlikely to occur in a category such as "clothes", where consumers are likely to complement each choice with other products from the same category.

[27]The average share is calculated as the mean of each product's consideration share. The share of consumers who consider a product j is given by: $(1/I)\left(\sum_j \sum_{c_t \in G} \delta_{ijt}^{c_t} P_{it}(c_t)\right)$ where I is the total number of consumers.

[28]For the elasticities and counterfactuals reported below, we choose the specification where the saliency parameter has not been disaggregated for new and existing products (columns (1) and (2) of Table 11) We do this to examine the overall saliency effect, which simplifies our policy counterfactuals as well as the interpretation of the results.

corresponding to the top-selling product in the sub-category). Each entry reports the mean elasticity across the 90 consumers who shopped four travel bags. The variation in estimated elasticities (given by the ratio of the maximum to the minimum cross-price elasticity within a column) ranges from -2 to 0. This indicates that the model has overcome the restrictive form imposed by the multinomial logit model which produces proportional substitution elasticities (Nevo 2000). Most cross-elasticities are negative[29] indicating that recommendations generate negative externalities, contracting the market for other products. Some cross-elasticities are positive, indicating positive externalities. This could be a result of spillovers from the saliency effect through cross-recommendations as documented in section 5.5. Overall, the average change in sales for all products as a result of a change in saliency for any product j , is negative but much smaller than own-elasticity effects.

Table 12 also allows us to investigate whether observed increases in demand for more salient products are driven by more customers purchasing these products (extensive margin) or whether customers substitute salient products for products that were not made more salient within product categories (intensive margin). The estimated own- and cross-price elasticities suggest that effects at both extensive and intensive margins are at work. For most products, cross-elasticities are overall negative, which suggests some substitution between products that received a saliency shock and other products which indicates effects are driven by changes at the intensive margin. That said, a comparison between own- and cross-elasticities suggests that for 8 out of 10 products this substitution effect is outweighed by an increase in demand for the more salient product. In other words, the demand increase occurs also at the extensive margin, i.e., new customers purchasing a salient product within a given product category.

6.3 Counterfactuals

Based on our model estimates, next we simulate counterfactuals to assess how sales shares change when: (1) consumers have full, instead of limited, attention (Kim et al. 2010) and (2) only certain types of products receive recommendations. Lewis and Wang (2013) show theoretically that while recommendation systems help generate a positive surplus for the platform as a result of improved matches, the overall effect may not be Pareto improving, as market participation may shift in favor of some products against others. We investigate this issue empirically, by examining how different types of recommendations systems impact the sales share of popular vs. unpopular products, when consumers have limited attention.

For the full attention case, our model reduces to a standard multinomial logit (MNL) model; we thus use the estimated parameters from the MNL model to calculate sales shares when consumers have full attention. To assess the performance of different recommendation systems, we define the set of "popular products" as those with an observed sales share of more than 10% (products 1, 2 and 3). The remaining products (products 4 to 12) are categorized as "unpopular products." We then consider two variants of the recommendation system: i)

[29]Note, that the zero elasticities are a result of the fact that saliency is zero for these (column) products.

when only popular products are (always) recommended (products 1, 2 and 3) and ii) when only unpopular products are (always) recommended (products 4 to 12).

Figure 16 plots the results from the simulations. Each bar in the figure represents the percentage difference in sales share between when consumers have limited attention and when they have full attention. A negative value indicates that the share under limited information is lower than that under full attention. The x-axis orders products by their sales rank (1 being most popular and 12 being least popular).

The first result is that limited attention disproportionately harms top-selling products. Popular products suffer a loss in their sales share when consumers have limited attention under the existing recommendation system as indicated by the dark grey bars. Intuitively, popular products are chosen less when consumers have limited attention because they enter less frequently in their consideration sets. Under full attention however, they are always in a consumer's consideration set and are frequently chosen based on their superior underlying characteristics. As the actual recommendation system does not disproportionately highlight popular products (Table 11 reports that 10 out of 12 products in this sub-category are recommended by a new product at least once), their share under limited information is always lower relative to the full attention scenario.

Next, we consider a different recommendation system, whereby, all new products recommend only popular products, namely products 1, 2 and 3. Under this scenario (medium gray bars), we find the difference in sales share for popular products 1 and 2 is still negative but the loss that they suffer is much lower than when all products are equally recommended. Popular product 3 actually reverses the share differential and stands to gain by almost 2% under this system. In contrast, some unpopular products (7, 10, 11 and 12) register a negative difference as a result of only popular products being recommended.

Finally we plot how sales shares evolve under limited and full attention, if only unpopular products received recommendations from new products. The light grey bars show that this would, predictably, lead to an increase in the share of unpopular products (under limited attention) and cause the sales share difference to increase by almost about 1.5%. Unpopular products benefit from consumers having limited attention (relative to full attention) because popular products will be even more overlooked under this recommendation system. Popular products in contrast lose up to 4.5% of sales share under this setting.

7 Conclusion

In this paper we estimate the effect of product saliency, affected via recommendation sets, on user choice in online markets. We find a sharp and robust 6% increase in the sales of a product when it is recommended by a highly popular new product. This effect is however short-lived, lasting for approximately only four days. On average the daily increase in sales attributable to saliency is around 5%. We also find that products recommended in smaller sets experience larger effects of saliency as they have to compete less for user attention. Finally, our context allows us to build a robust counterfactual to verify our results. We exploit regional variation in recommendation sets whereby we compare regional sales of products that receive recommendations from the same product but over different regions. We find that products recommended by a new product in America see a 13% increase in their American-Europe sales differential over the 4 event days, compared to similar products recommended exclusively in Europe by the same new product. Our structural analysis confirms that the reduced form effects of saliency on sales are the result of saliency affecting the set of products considered by a consumer. Once we condition on the effect of saliency on the consideration set, saliency has no effect on choice.

Our analysis sheds light on consumer choice in an environment with low search costs but very large choice sets. Our evidence rejects the traditional revealed preference assumption and suggests that consumers make choices consistent with revealed attention. In other words, we find that consumers make choices under limited attention despite having easy access to information. This results in a search friction as consumers do not consider all products available on a given online platform. However, we also show that online platforms operating in these environments can alleviate this friction by helping refine user search by offering product recommendations. We show that such recommendations can temporarily affect the sales and appeal for products by shaping and expanding consumers' consideration sets. Indeed, our counterfactual analysis suggests that different recommendation systems can have large effects on the demand for a given product if consumer choice is subject to limited attention.

References

ABADIE, A., A. DIAMOND, AND J. HAINMUELLER (2010): "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program," *Journal of the American Statistical Association*, 105.

ANDREWS, R. L. AND T. SRINIVASAN (1995): "Studying consideration effects in empirical choice models using scanner panel data," *Journal of Marketing Research*, 30–41.

BAŞAR, G. AND C. BHAT (2004): "A parameterized consideration set model for airport choice: an application to the San Francisco Bay area," *Transportation Research Part B: Methodological*, 38, 889–904.

BAYE, M. R., J. R. J. GATTI, P. KATTUMAN, AND J. MORGAN (2009): "Clicks, discontinuities, and firm demand online," *Journal of Economics & Management Strategy*, 18, 935–975.

CAI, H., Y. CHEN, AND H. FANG (2009): "Observational Learning: Evidence from a Randomized Natural Field Experiment," *The American Economic Review*, 864–882.

CAMERON, A. C. AND P. K. TRIVEDI (2013): *Regression analysis of count data*, 53, Cambridge university press.

CAPLIN, A., M. DEAN, AND D. MARTIN (2011): "Search and satisficing," *The American Economic Review*, 2899–2922.

CARARE, O. (2012): "The Impact Of Bestseller Rank On Demand: Evidence From The App Market," *International Economic Review*, 53, 717–742.

CHETTY, R., A. LOONEY, AND K. KROFT (2009): "Salience and Taxation: Theory and Evidence," *The American Economic Review*, 99, 1145.

CHIANG, J., S. CHIB, AND C. NARASIMHAN (1998): "Markov chain Monte Carlo and models of consideration set and parameter heterogeneity," *Journal of Econometrics*, 89, 223–248.

CONLON, C. T. AND J. H. MORTIMER (2013): "Demand Estimation Under Incomplete Product Availability," *American Economic Journal: Microeconomics*, 5, 1–30.

CRAWFORD, G. S., R. GRIFFITH, AND A. IARIA (2015): "Estimating Demand Parameters with Unobserved Choice Sets," *mimeo*.

DE LOS SANTOS, B., A. HORTAÇSU, AND M. R. WILDENBEEST (2012): "Testing models of consumer search using data on web browsing and purchasing behavior," *The American Economic Review*, 102, 2955–2980.

DELLAVIGNA, S. (2009): "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, 47, 315–72.

DINERSTEIN, M., L. EINAV, J. LEVIN, AND N. SUNDARESAN (2014): "Consumer price search and platform design in internet commerce," Tech. rep., National Bureau of Economic Research.

Draganska, M. and D. Klapper (2011): "Choice set heterogeneity and the role of advertising: An analysis with micro and macro data," *Journal of Marketing Research*, 48, 653–669.

Eliaz, K. and R. Spiegler (2011): "Consideration sets and competitive marketing," *The Review of Economic Studies*, 78, 235–262.

Fleder, D. M. and K. Hosanagar (2007): "Recommender systems and their impact on sales diversity," in *Proceedings of the 8th ACM conference on Electronic commerce*, ACM, 192–199.

Goeree, M. S. (2008): "Limited information and advertising in the US personal computer industry," *Econometrica*, 76, 1017–1074.

Haan, M. A. and J. L. Moraga-González (2011): "Advertising for attention in a consumer search model," *The Economic Journal*, 121, 552–579.

Hauser, J. R. and B. Wernerfelt (1990): "An evaluation cost model of consideration sets," *Journal of consumer research*, 393–408.

Honka, E., A. Hortaçsu, and M. A. Vitorino (2014): "Advertising, Consumer Awareness and Choice: Evidence from the US Banking Industry," *Unpublished manuscript*.

Huang, J.-H. and Y.-F. Chen (2006): "Herding in online product choice," *Psychology & Marketing*, 23, 413–428.

Kawaguchi, K., K. Uetake, and Y. Watanabe (2014): "Identifying Consumer Inattention: A Product-Availability Approach," *Available at SSRN*.

Kim, J. B., P. Albuquerque, and B. J. Bronnenberg (2010): "Online demand under limited consumer search," *Marketing science*, 29, 1001–1023.

Lewis, G. and A. Wang (2013): "Who benefits from improved search in platform markets?" *mimeo*.

Manski, C. F. (1977): "The structure of random utility models," *Theory and decision*, 8, 229–254.

Manzini, P. and M. Mariotti (2014): "Stochastic choice and consideration sets," *Econometrica*, 82, 1153–1176.

Masatlioglu, Y., D. Nakajima, and E. Y. Ozbay (2012): "Revealed attention," *The American Economic Review*, 102, 2183–2205.

Mozer, M. C. and M. Sitton (1998): "Computational modeling of spatial attention," *Attention*, 9, 341–393.

Nevo, A. (2000): "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand," *Journal of Economics & Management Strategy*, 9, 513–548.

Roberts, J. H. and J. M. Lattin (1991): "Development and testing of a model of consideration set composition," *Journal of Marketing Research*, 429–440.

Schafer, J. B., D. Frankowski, J. Herlocker, and S. Sen (2007): "Collaborative filtering recommender systems," in *The adaptive web*, Springer, 291–324.

Senecal, S. and J. Nantel (2004): "The influence of online product recommendations on consumers' online choices," *Journal of retailing*, 80, 159–169.

Sims, C. A. (2003): "Implications of rational inattention," *Journal of monetary Economics*, 50, 665–690.

Smith, M. D. and E. Brynjolfsson (2001): "Consumer decision-making at an Internet shopbot: Brand still matters," *The Journal of Industrial Economics*, 49, 541–558.

Sorensen, A. T. (2007): "Bestseller Lists And Product Variety," *The Journal of Industrial Economics*, 55, 715–738.

Swait, J. and M. Ben-Akiva (1987): "Incorporating random constraints in discrete models of choice set generation," *Transportation Research Part B: Methodological*, 21, 91–102.

Tucker, C. and J. Zhang (2011): "How does popularity information affect choices? A field experiment," *Management Science*, 57, 828–842.

Van Nierop, E., B. Bronnenberg, R. Paap, M. Wedel, and P. H. Franses (2010): "Retrieving unobserved consideration sets from household panel data," *Journal of Marketing Research*, 47, 63–74.

Wooldridge, J. M. (1999): "Distribution-free estimation of some nonlinear panel data models," *Journal of Econometrics*, 90, 77–97.

Zajonc, R. B. (1968): "Attitudinal effects of mere exposure," *Journal of Personality and Social Psychology Monographs*, 9.

Figure 1: Recommendation sets

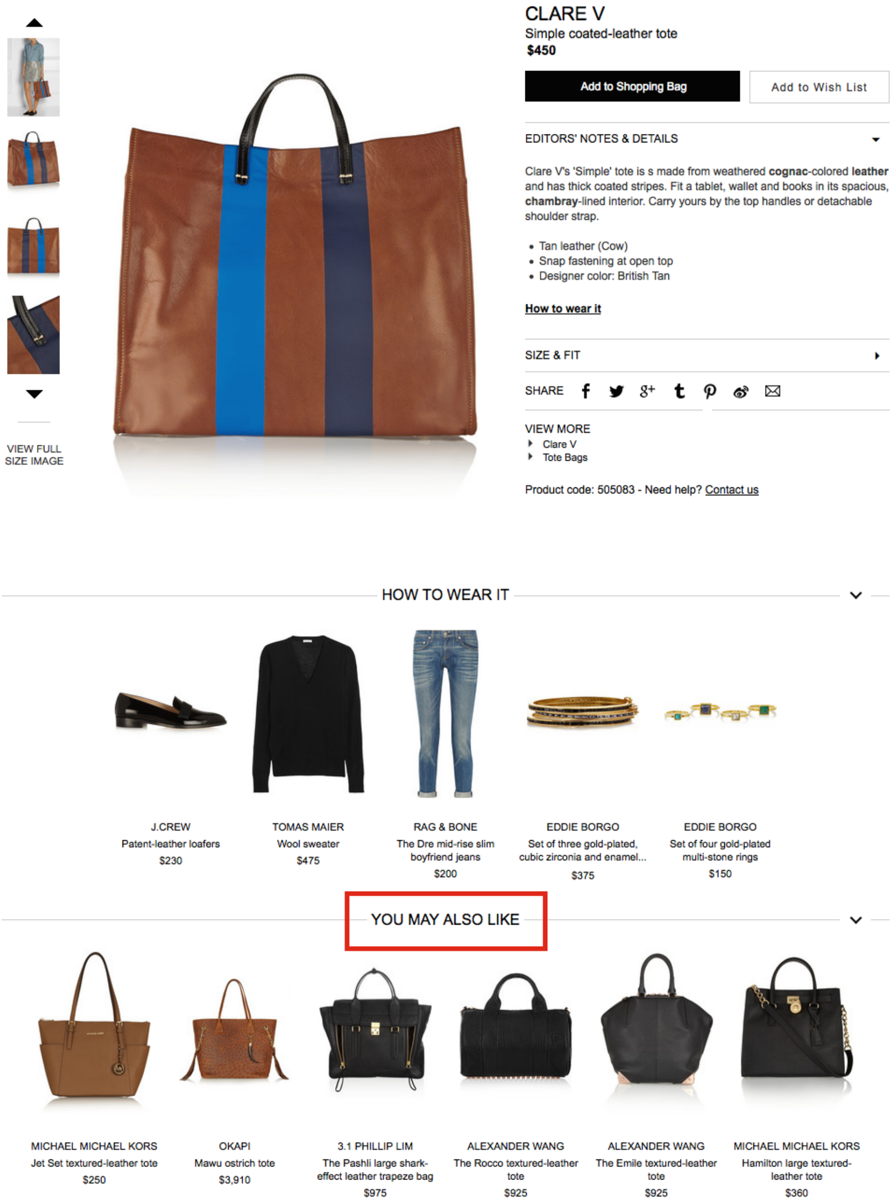

Figure 2: Recommendation sets U.S. vs. Europe

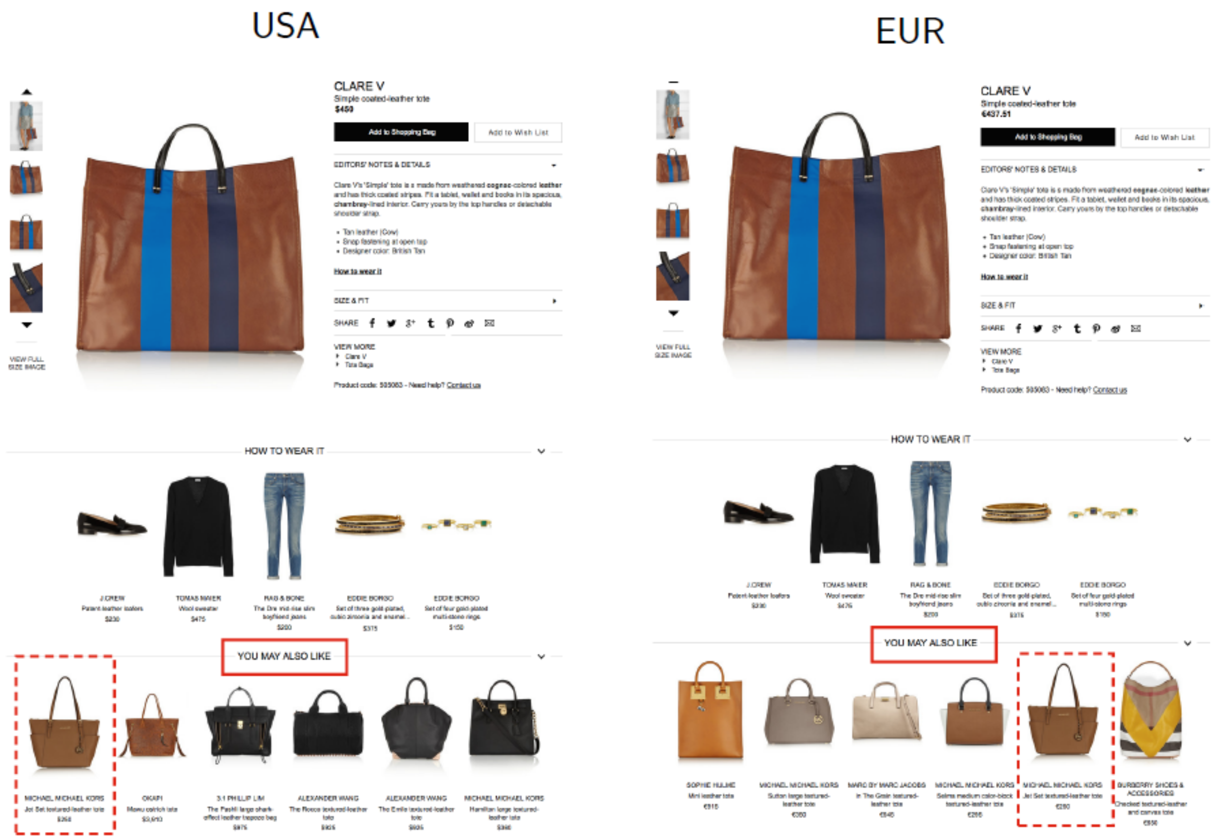

Figure 3: Differences in shopping bag additions between existing products recommended by new products and all other products

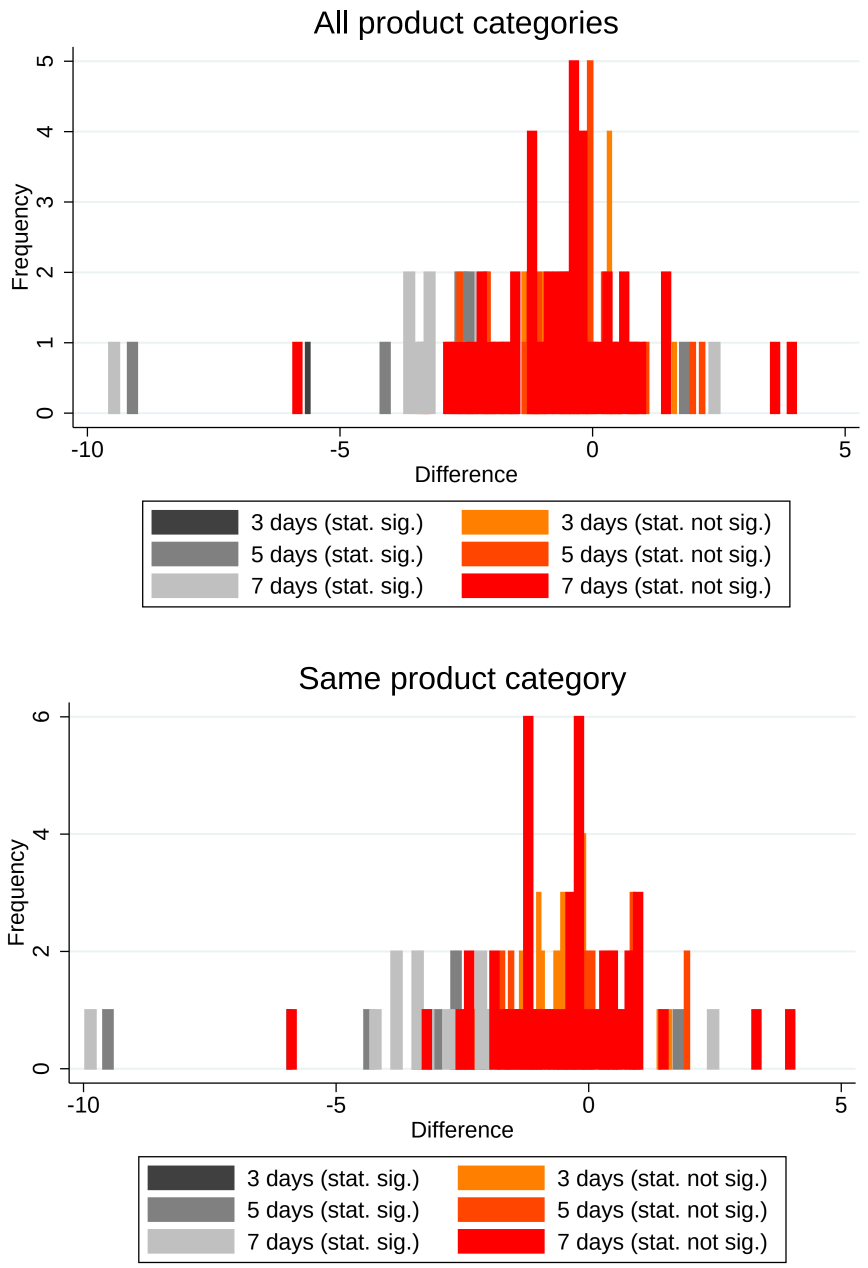

Note: *difference* computed as difference between the number of shopping bag additions for an existing product during 3, 5, or 7 days prior to the recommendation by a new product and the number of shopping bag additions for other existing products that are not recommended by a new product during the same time period. Each bar corresponds to the arrival date of a new product. *Same product category* means we only consider existing products that are recommended by new products that are in the same product category as the existing product.

Table 1: Descriptive statistics: Net-a-Porter ‘Live’

Variable	Category	Existing product not recommended by new product		Existing product recommended by new product		t-test	New product not recommended by new product		New product recommended by new product		t-test
		Mean	Std. Dev.	Mean	Std. Dev.		Mean	Std. Dev.	Mean	Std. Dev.	
# Shopping bag	All	0.385	1.062	0.555	1.313	-0.170***	0.278	1.284	0.539	1.879	-0.261***
	Accessories	0.217	0.871	0.333	1.109	-0.116***	0.203	1.105	0.395	1.507	-0.191***
	Bags	0.352	1.088	0.566	1.354	-0.213***	0.252	1.172	0.460	1.522	-0.208***
	Beauty	0.499	1.194	0.585	1.290	-0.085***	0.474	1.597	0.549	2.198	-0.075***
	Clothing	0.359	0.948	0.532	1.217	-0.173***	0.269	1.306	0.548	1.894	-0.279***
	Lingerie	0.529	1.140	0.706	1.280	-0.176***	0.333	1.420	0.580	1.960	-0.246***
	Shoes	0.611	1.494	0.900	1.831	-0.289***	0.387	1.412	0.763	2.370	-0.376***
# Wishlist	All	0.179	0.677	0.246	0.674	-0.067***	0.188	1.080	0.395	1.507	-0.121***
	Accessories	0.101	0.405	0.152	0.626	-0.050***	0.139	0.834	0.460	1.522	-0.107***
	Bags	0.236	1.388	0.248	0.662	-0.012**	0.210	1.220	0.549	2.198	-0.091***
	Beauty	0.143	0.484	0.168	0.527	-0.024***	0.177	0.860	0.548	1.894	-0.041***
	Clothing	0.182	0.533	0.251	0.648	-0.069***	0.192	1.137	0.580	1.960	-0.128***
	Lingerie	0.157	0.461	0.205	0.527	-0.048***	0.153	0.890	0.763	2.370	-0.077***
	Shoes	0.281	0.720	0.427	0.921	-0.146***	0.260	1.198	0.310	1.382	-0.163***
# Users		38,437	67,620	46,561	75,658	-8,124***	63,110	164,521	91,687	198,239	-28,576***

Transaction-level data aggregated at the day-level. # *Shopping bag* means number of shopping bag additions; # *Wishlist* means number of wishlist additions.

Table 2: Descriptive statistics: recommended and recommending products

	Same category			Same designer			Price difference†		
	Mean	Std. Dev.	Median	Mean	Std. Dev.	Median	Mean	Std. Dev.	Median
All	0.989	0.074	1	0.379	0.391	0.250	$9.210	$1,111	$-1.666
Accessories	0.974	0.118	1	0.388	0.376	0.285	$-0.117	$1,519	$-23.625
Bags	0.967	0.135	1	0.576	0.397	0.666	$32.987	$898	$0
Beauty	0.995	0.058	1	0.665	0.397	1	$1.651	$47	$0
Clothing	0.996	0.041	1	0.280	0.350	0.111	$9.116	$1,205	$0
Lingerie	0.994	0.053	1	0.695	0.385	1	$-19.873	$282	$0
Shoes	0.986	0.077	1	0.441	0.392	0.400	$24.092	$370	$-5.000

† Computed as price of recommended product minus price of recommending product.

Table 3: Descriptive statistics: average prices of products recommended by new products and non-recommended products

Price (in US$)	Recommended	Not recommended	Difference
All	1,075.34	1,036.86	38.47
Accessories	989.08	1,045.29	-56.21
Bags	1,666.98	1,732.24	-65.25
Beauty	71.65	78.99	-7.34
Clothing	1,250.59	1,281.34	-30.75
Lingerie	258.28	222.23	-36.04**
Shoes	800.14	873.05	-72.91**

** Difference significant at 5%.

Table 4: Descriptive statistics: product recommendations by new products

# Products	Recommended by new product		If recommended # recommendations		
	# products	% Total	Mean	SD	Max
By product recommended by new product					
All	6,693	43.46	1.50	0.91	23
Accessories	1,197	52.61	1.49	0.79	8
Bags	590	50.77	1.49	0.75	6
Beauty	310	31.47	1.55	1.39	9
Clothing	3,523	42.24	1.53	0.84	8
Lingerie	240	27.71	1.98	2.30	23
Shoes	833	47.06	1.45	0.79	8
By product recommended by new product and time					
All	373.57	2.95	1.27	0.68	23
Accessories	64.19	3.56	1.28	0.62	8
Bags	32.63	3.57	1.27	0.58	6
Beauty	6.98	0.75	1.44	1.26	9
Clothing	224.68	3.34	1.26	0.61	8
Lingerie	6.37	0.79	1.84	2.09	23
Shoes	38.68	2.62	1.24	0.61	8

The upper part of this table reports statistics when we consider whether an existing product was recommended by a new product during the entire period of observation. The lower part of the table, in contrast, reports statistics when we consider whether an existing product was recommended by a new product on a given day during the period of observation. The # of products recommended by new products in the lower part of the table are therefore the average number of existing products recommended by new products on any given day.

Figure 4: Saliency Shock: New Product

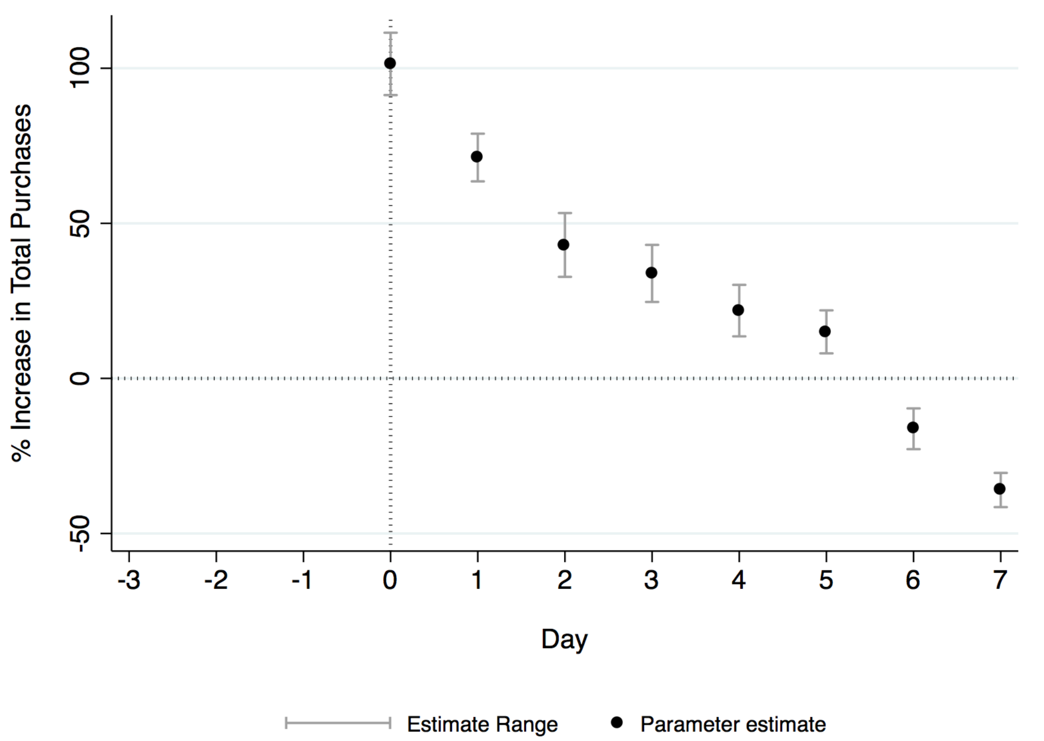

This figure reports coefficient estimates (with 95% confidence intervals) of the effect of entry – being a new product – on the total number of **shopping bag** additions (per day). The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *New Product* is a dummy variable taking the value one is the product was introduced in the catalogue on a given day. The regression specifications controls for time fixed effects, controls for day of the week and weekend.

Table 5: Effects of Saliency on Total Purchases

	(1)	(2)	(3)	(4)
Saliency	0.059***	0.057***	0.055***	0.054***
	(0.013)	(0.014)	(0.016)	(0.016)
Forward Lag of Saliency		-0.027	-0.027	-0.029
		(0.017)	(0.017)	(0.018)
New Product	0.967***	0.981***	0.980***	0.954***
	(0.040)	(0.040)	(0.041)	(0.041)
Saliency × New Product			0.003	0.004
			(0.022)	(0.022)
2 Week Lag of New Product				0.069**
				(0.027)
Time and Day F.E	Yes	Yes	Yes	Yes
Observations	986214	969368	969368	969368

This table reports results on the effect of the saliency shock on the total number of **shopping bag** additions (per day). The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *New Product* is a dummy variable taking the value one is the product was introduced in the catalogue on a given day. *Saliency* is defined as the total number of new products that recommend the target product at any given point of time. All specifications control for time fixed effects, controls for day of the week and weekend. Standard errors clustered by product are reported in parentheses. * indicates significance at 10%; ** at 5%; *** at 1%.

Table 6: Effects of Saliency on Total Wishlist

	(1)	(2)	(3)	(4)
Saliency	0.056***	0.054***	0.063***	0.060***
	(0.011)	(0.011)	(0.017)	(0.017)
Forward Lag of Saliency		-0.005	-0.004	-0.007
		(0.018)	(0.018)	(0.018)
New Product	1.181***	1.187***	1.192***	1.130***
	(0.029)	(0.029)	(0.030)	(0.029)
Saliency × New Product			-0.014	-0.010
			(0.021)	(0.021)
2 Week Lag of New Product				0.210***
				(0.020)
Time and Day F.E	Yes	Yes	Yes	Yes
Observations	932539	915993	915993	915993

This table reports results on the effect of the saliency shock on the total number of **wish-list additions** (per day). The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *New Product* is a dummy variable taking the value one is the product was introduced in the catalogue on a given day. *Saliency* is defined as the total number of new products that recommend the target product at any given point of time. All specifications control for time fixed effects, controls for day of the week and weekend. Standard errors clustered by product are reported in parentheses. * indicates significance at 10%; ** at 5%; *** at 1%.

Figure 5: Staggered Effects of Saliency Shocks

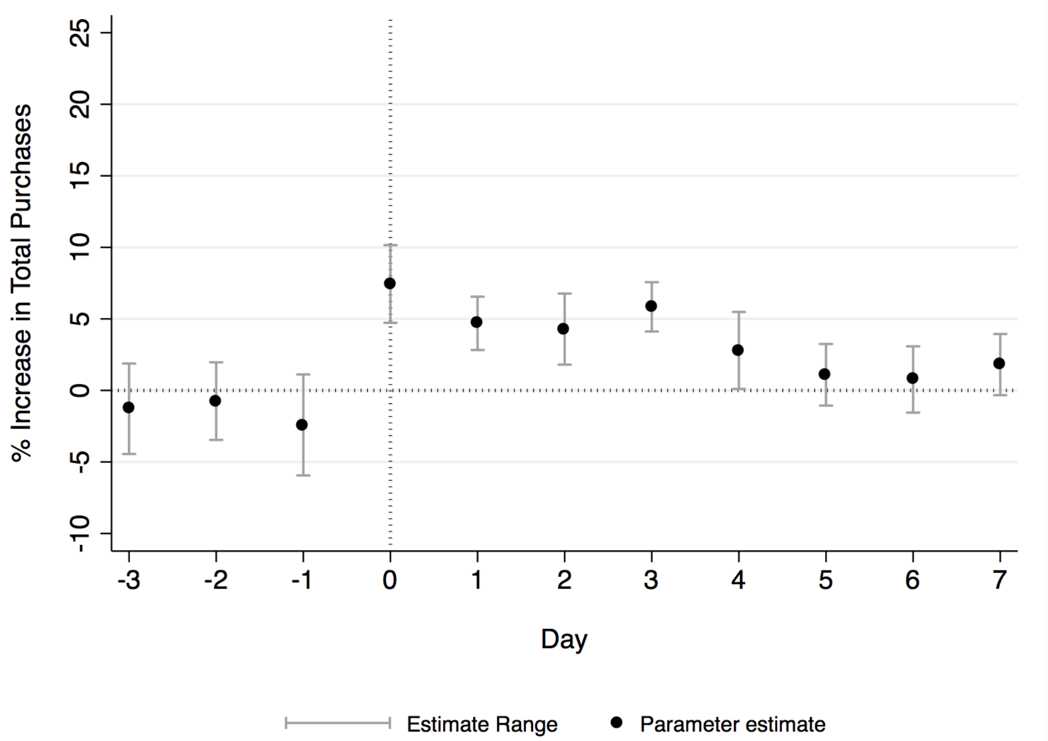

This figure reports coefficient estimates (with 95% confidence intervals) of the effect of saliency on the total number of **shopping bag** additions (per day). The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *Saliency* is defined as the total number of new products that recommend the target product at any given point of time. The regression specifications controls for the entry of a new product, time fixed effects, controls for day of the week and weekend.

Figure 6: Staggered Effects of Saliency Shocks

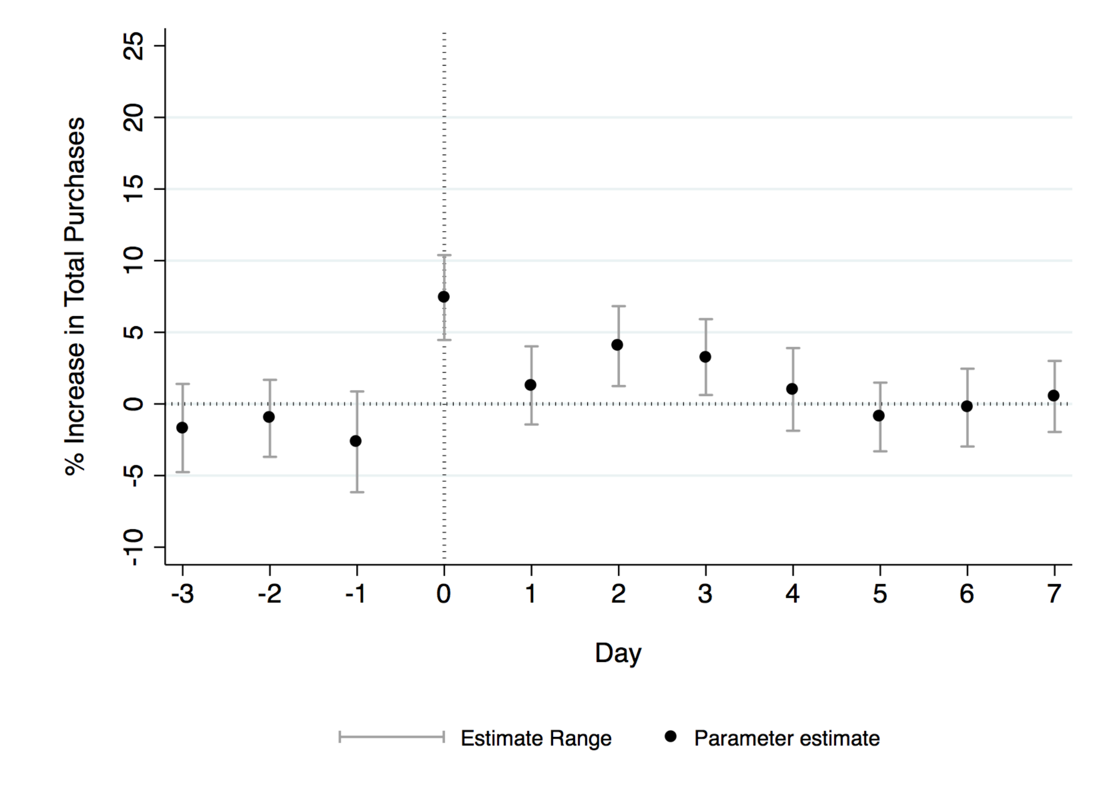

(a) **Saliency**

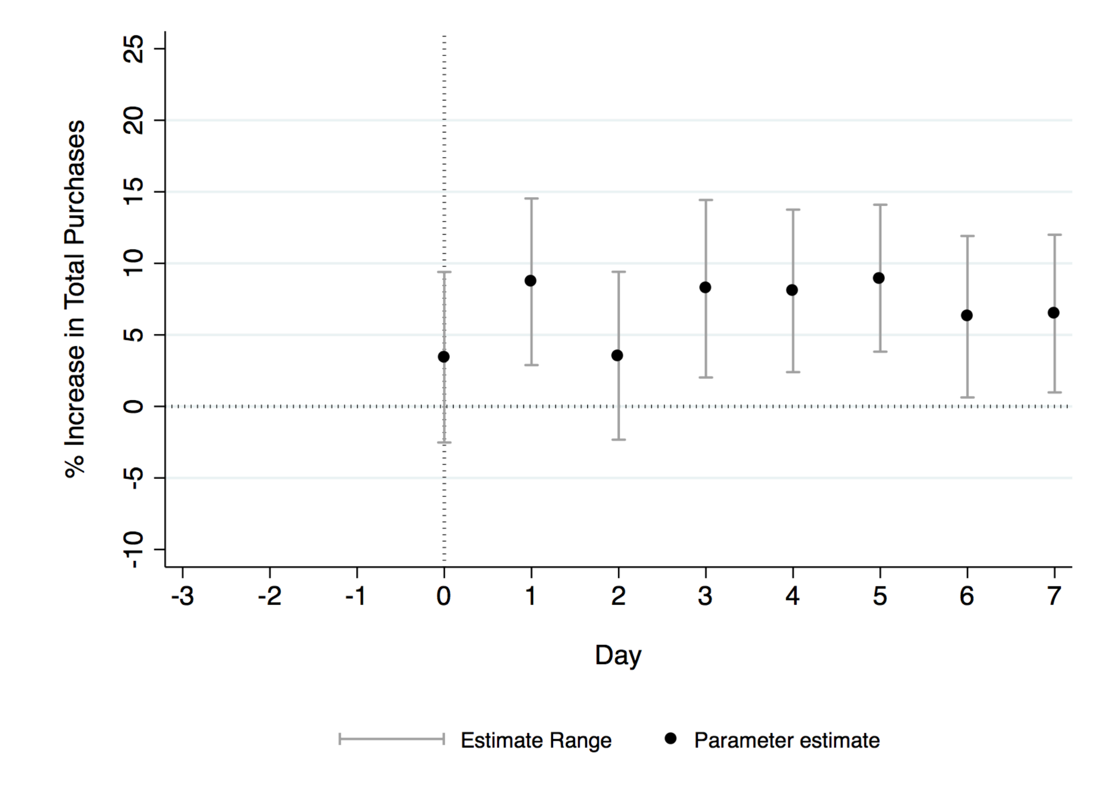

(b) **New Product** × **Saliency**

These figures reports coefficient estimates (with 95% confidence intervals) of the effect of saliency on the total number of **shopping bag** additions (per day). Figure (a) reports the effect of saliency for existing products while figure (b) reports the effect of saliency for new products. The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *Saliency* is defined as the total number of new products that recommend the target product at any given point of time. The regression specifications controls for the entry of a new product, time fixed effects, controls for day of the week and weekend.

Figure 7: Saliency Shocks: : Controlling for Lagged Effect of New Product

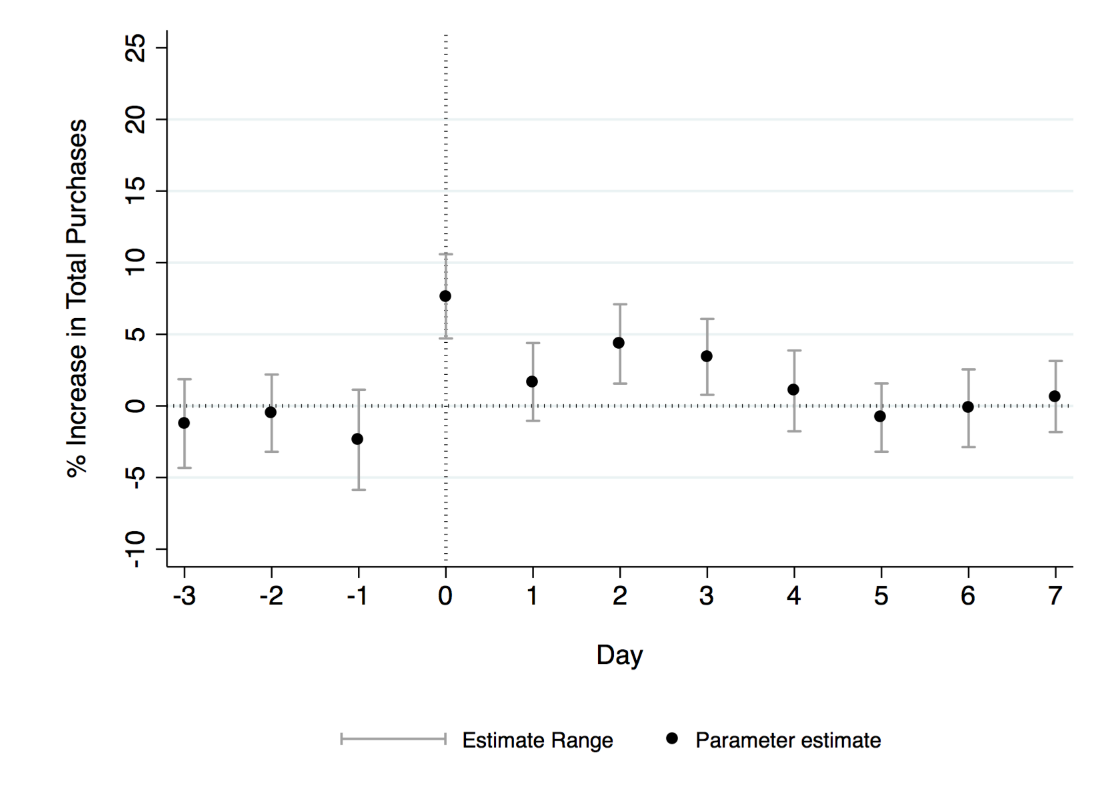

(a) **Saliency**

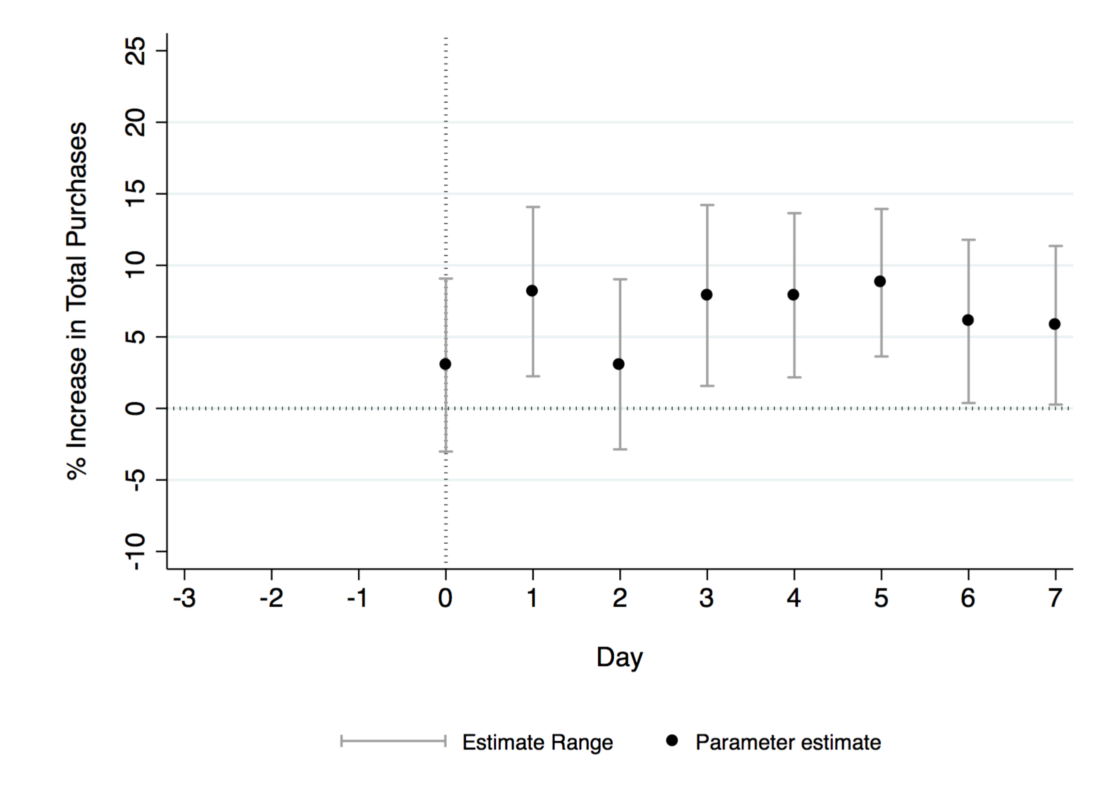

(b) **New Product × Saliency**

These figures reports coefficient estimates (with 95% confidence intervals) of the effect of saliency on the total number of **shopping bag** additions (per day) controlling for lagged entry effects, i.e, we include the two-week lag of whether a product was new in the specification. Figure (a) reports the effect of saliency for existing products while figure (b) reports the effect of saliency for new products. The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *Saliency* is defined as the total number of new products that recommend the target product at any given point of time. The regression specifications controls for the entry of a new product, time fixed effects, controls for day of the week and weekend.

Figure 8: Saliency Shocks: Controlling for Differential Anticipation Effects

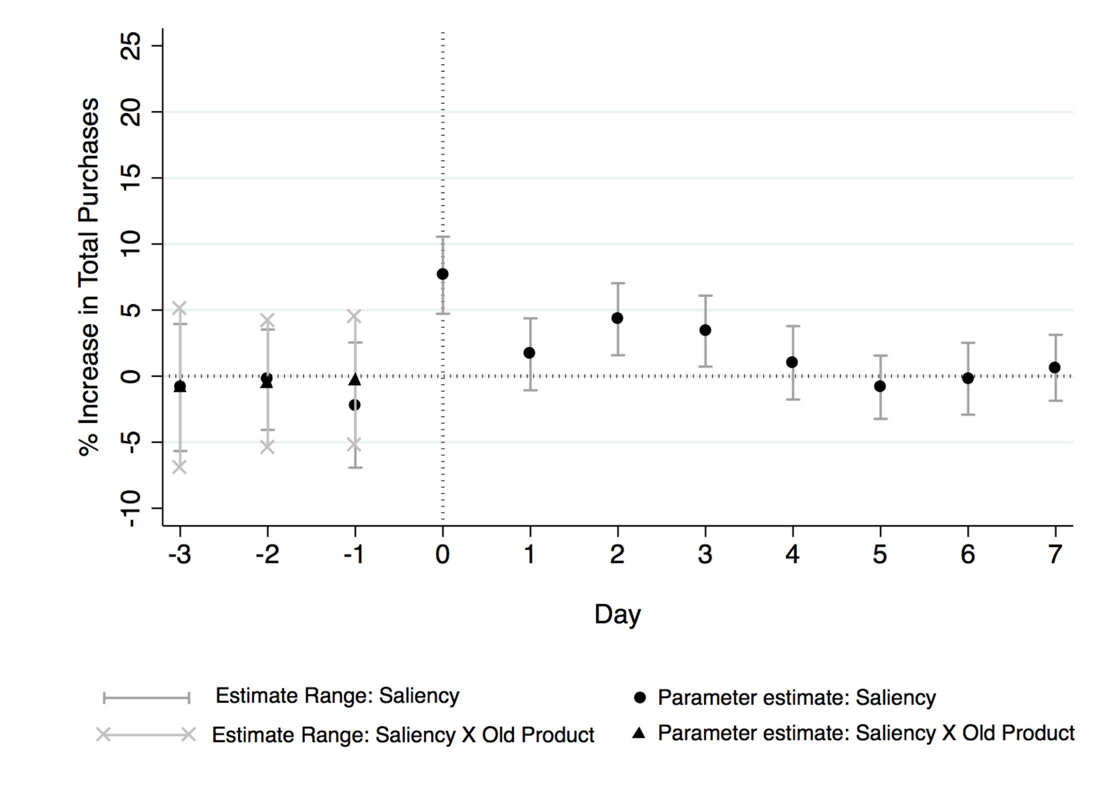

(a) **Saliency**

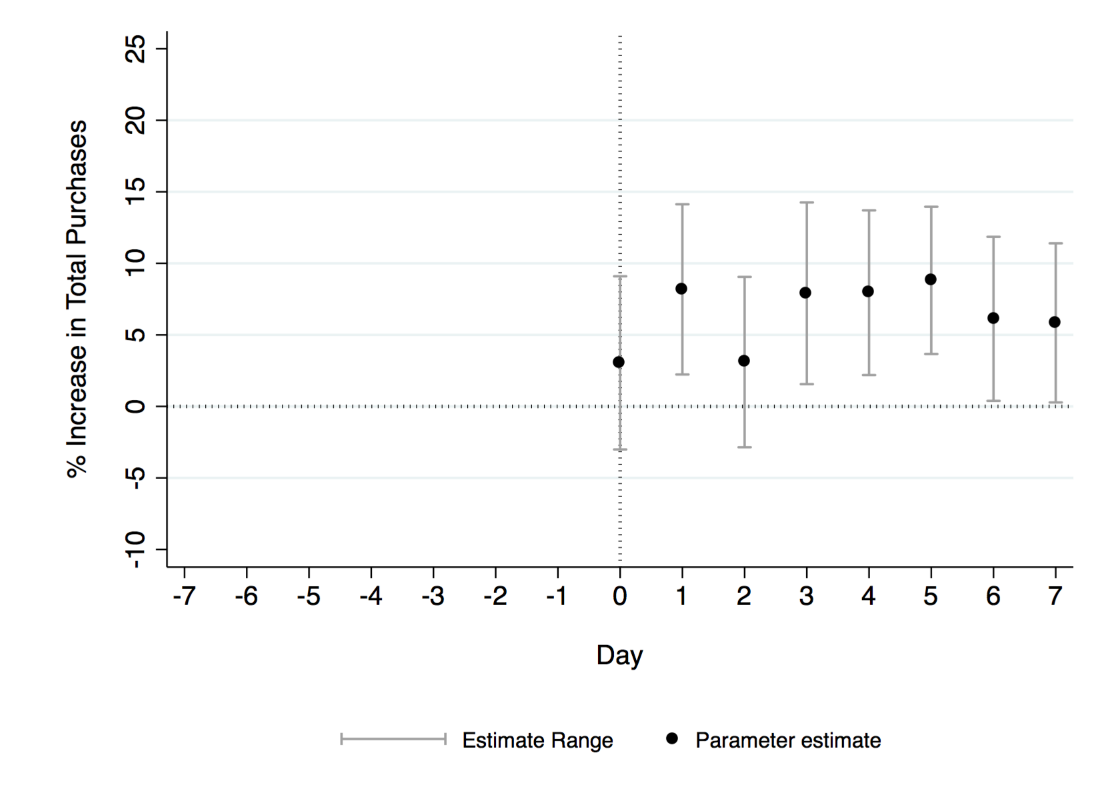

(b) **New Product × Saliency**

These figures reports coefficient estimates (with 95% confidence intervals) of the effect of saliency on the total number of **shopping bag** additions (per day) controlling for differential anticipation effects, i.e, we split and include the anticipation effect (forward lags of saliency) between products that received a prior saliency shock and those that did not. Figure (a) reports the effect of saliency for existing products while figure (b) reports the effect of saliency for new products. The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *Saliency* is defined as the total number of new products that recommend the target product at any given point of time. The regression specifications controls for the entry of a new product, time fixed effects, controls for day of the week and weekend.

Figure 9: Saliency Shocks: Attention Effects

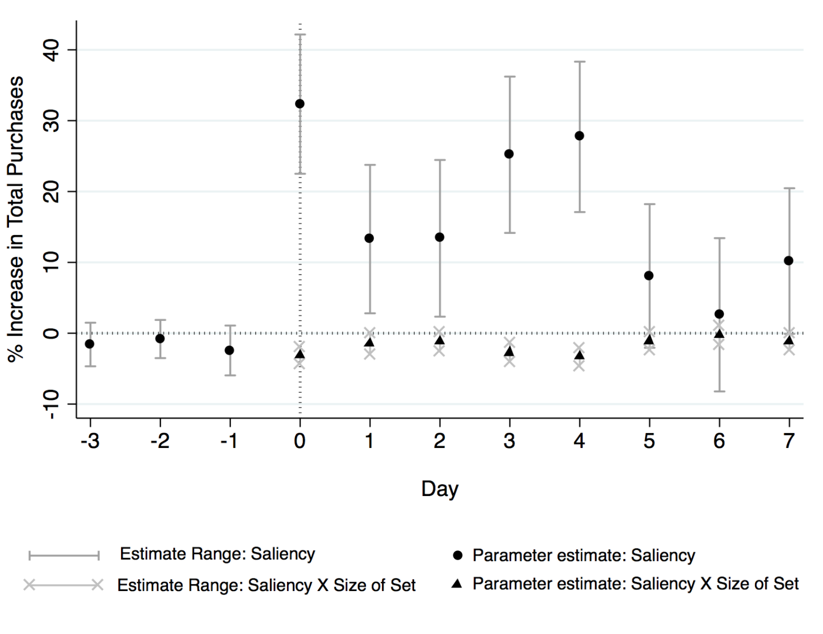

This figure reports coefficient estimates (with 95% confidence intervals) of the effect of saliency and its interaction with 'attention' (size of set) on the total number of **shopping bag** additions (per day). The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *Saliency* is defined as the total number of new products that recommend the target product at any given point of time. Our proxy for attention, the *size of the set* is defined as the average size of new product recommendation sets that include the target product. The regression specifications controls for the entry of a new product, time fixed effects, controls for day of the week and weekend.

Figure 10: Saliency Shocks: Price Effects

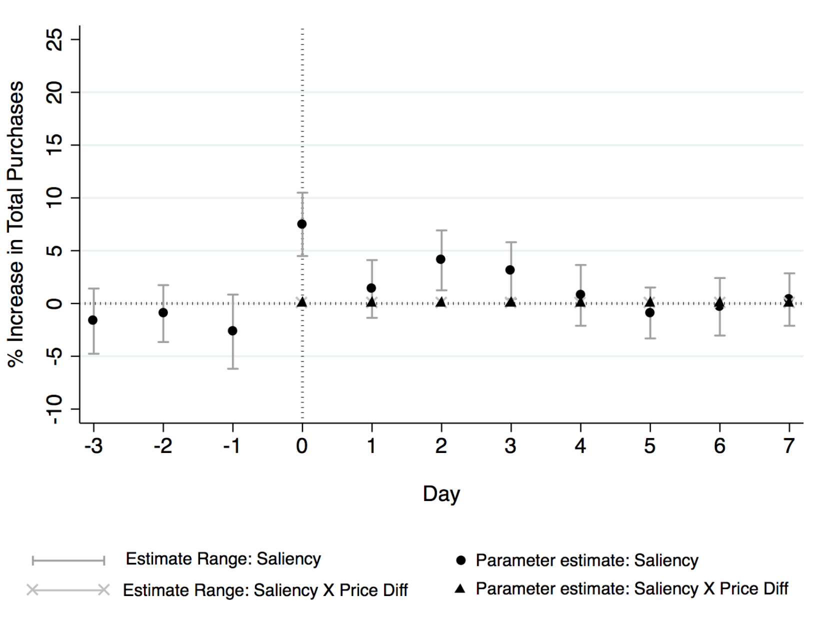

This figure reports coefficient estimates (with 95% confidence intervals) of the effect of saliency and its interaction with the price of the product on the total number of **shopping bag** additions (per day). The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *Saliency* is defined as the total number of new products that recommend the target product at any given point of time. *Price* is the retail price of the product in US dollars. The regression specifications controls for the entry of a new product, time fixed effects, controls for day of the week and weekend.

Figure 11: Saliency Shocks: Wishlist with Controls for Lagged New Product Effects

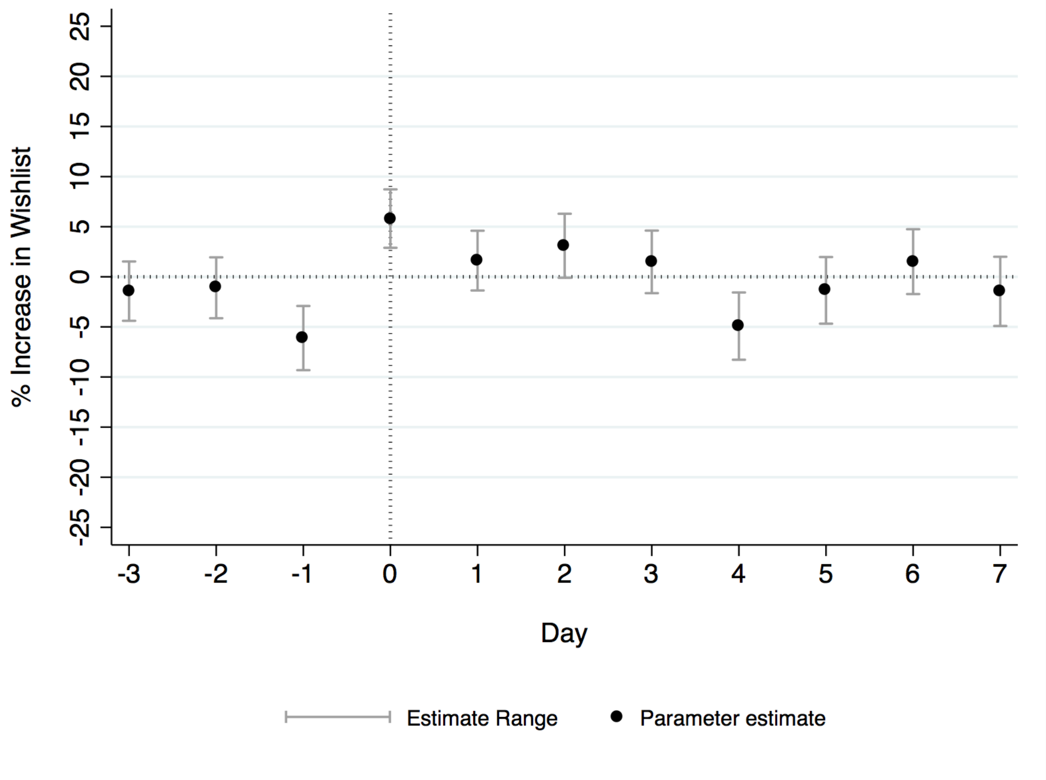

(a) **Saliency**

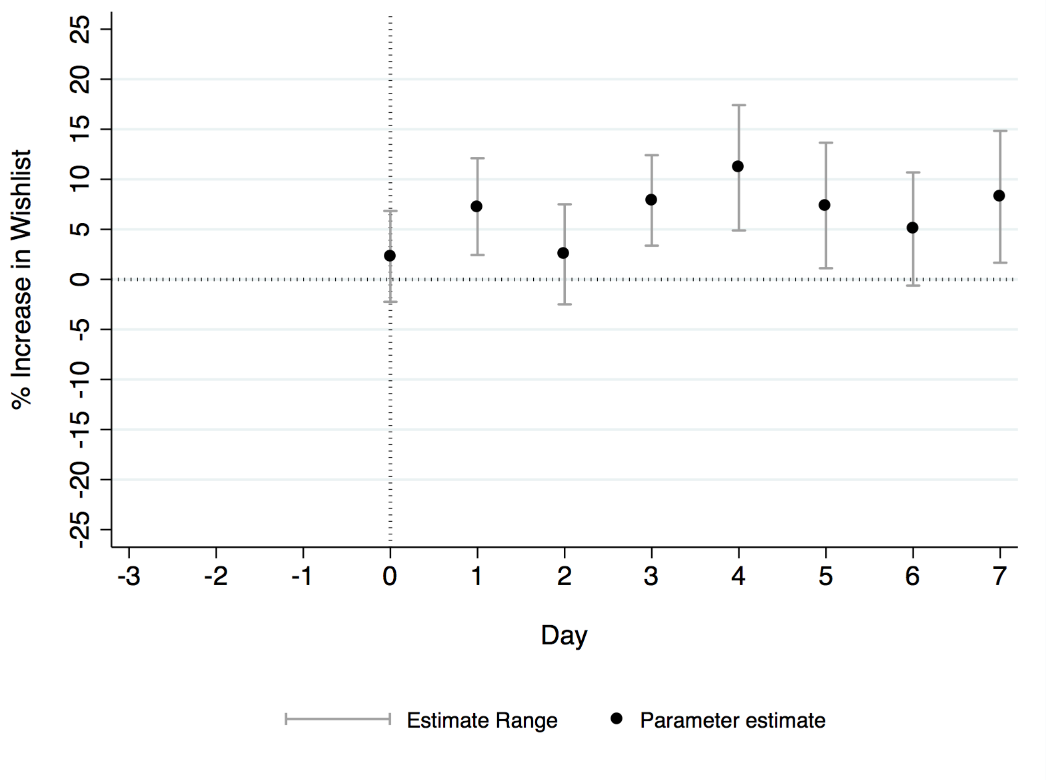

(b) **New Product** × **Saliency**

These figures reports coefficient estimates (with 95% confidence intervals) of the effect of saliency on the total number of **wish-list** additions (per day) controlling for lagged entry effects, i.e, we include the two-week lag of whether a product was new in the specification. Figure (a) reports the effect of saliency for existing products while figure (b) reports the effect of saliency for new products. The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *Saliency* is defined as the total number of new products that recommend the target product at any given point of time. The regression specifications controls for the entry of a new product, time fixed effects, controls for day of the week and weekend.

Figure 12: Saliency Shocks: Wishlist with Control for Differential Anticipation Effects

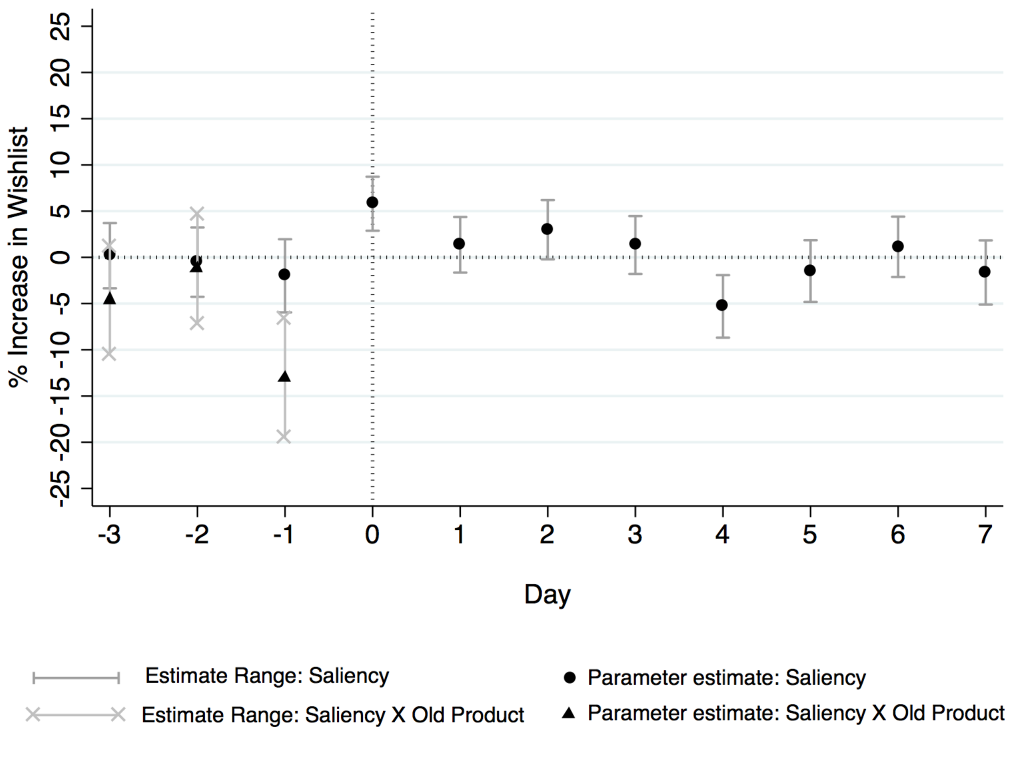

(a) **Saliency**

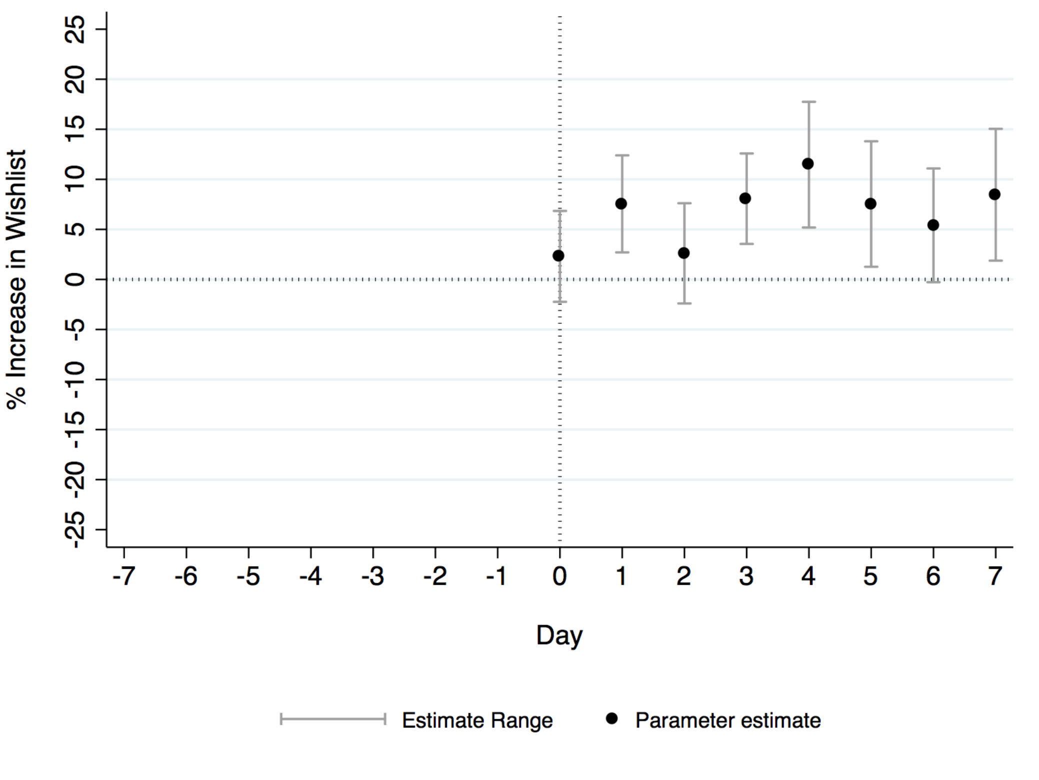

(b) **New Product × Saliency**

These figures reports coefficient estimates (with 95% confidence intervals) of the effect of saliency on the total number of **wish-list** additions (per day) controlling for differential anticipation effects, i.e, we split and include the anticipation effect (forward lags of saliency) between products that received a prior saliency shock and those that did not. Figure (a) reports the effect of saliency for existing products while figure (b) reports the effect of saliency for new products. The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *Saliency* is defined as the total number of new products that recommend the target product at any given point of time. The regression specifications controls for the entry of a new product, time fixed effects, controls for day of the week and weekend.

Figure 13: Staggered Effects of Saliency Shocks for Complementary Products

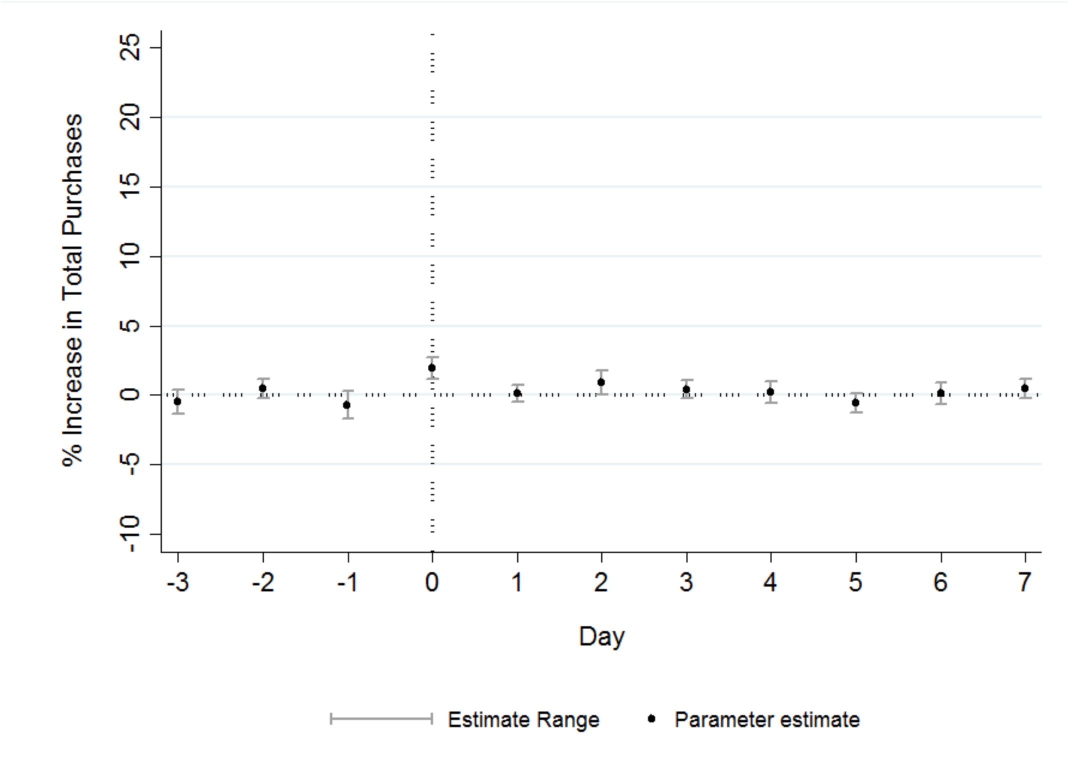

(a) **Saliency**

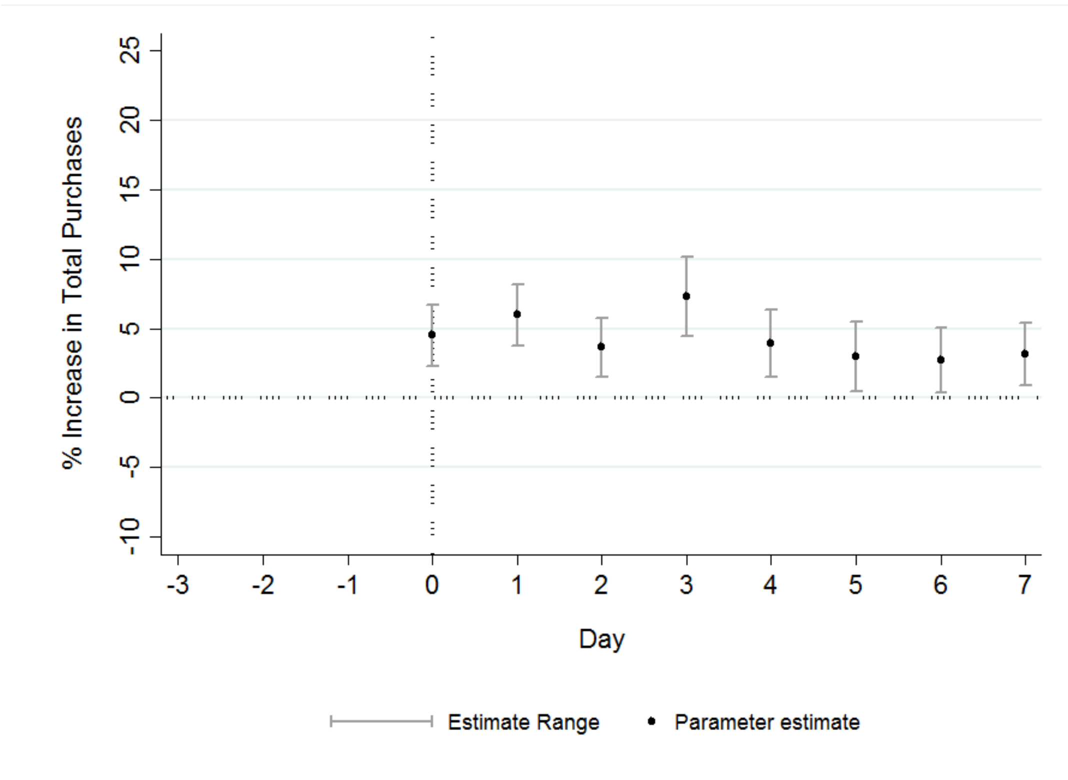

(b) **New Product** × **Saliency**

These figures reports coefficient estimates (with 95% confidence intervals) of the effect of saliency on the total number of **shopping bag** additions (per day). Figure (a) reports the effect of saliency for existing products while figure (b) reports the effect of saliency for new products. The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *Saliency* is defined as the total number of new products that recommend the target product **as complimentary** at any given point of time under the heading. The regression specifications controls for the entry of a new product, time fixed effects, controls for day of the week and weekend.

Table 7: Spillover Effects of Saliency

	-3	**-2**	**-1**	**Event day**	**+1**	**+2**	**+4**	**+5**	**+6**	**+7**	**+8**
Without spillovers:											
Direct saliency effect	-0.017	-0.01	-0.026	0.074***	0.013	0.040**	0.033**	0.01	-0.009	-0.003	0.005
	(0.019)	(0.016)	(0.021)	(0.018)	(0.017)	(0.017)	(0.016)	(0.018)	(0.015)	(0.017)	(0.015)
With spillovers:											
Direct saliency effect (degree 1)	-0.001	0.003	-0.018	0.082***	0.023	0.043**	0.038**	0.009	-0.009	0.01	0.008
	(0.019)	(0.017)	(0.022)	(0.018)	(0.017)	(0.017)	(0.016)	(0.018)	(0.015)	(0.017)	(0.016)
Indirect saliency effect (degree 2)	0.023***	0.023***	0.011*	0.019***	0.024***	0.013**	0.019***	0.010*	0.014**	0.029***	0.015***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.005)
Indirect saliency effect (degree 3)	0.006*	0.005	0.003	-0.002	0.002	-0.001	0.0001	-0.005	-0.003	0.008**	0.002
	(0.003)	(0.003)	(0.004)	(0.004)	(0.003)	(0.003)	(0.004)	(0.003)	(0.004)	(0.003)	(0.003)

This table reports results on the spill-over effects of the saliency shock on the total number of **shopping bag additions** (per day). The sample consists of all daily, product-level, transactions carried out between May 20, 2014 to July 29, 2014. As a result of the thrice-weekly introduction of new products, our database represents an unbalanced panel of products over time. *Direct Saliency (degree 1)* is defined as the total number of new products that recommend the target product at any given point of time. Products that are recommended by degree 1 products have a two degree separation between themselves and the new product and are identified by the variable *indirect Saliency (degree 2)*. Products that are recommended by degree 2 products have a three degree separation between themselves and the new product and are identified by the variable *indirect Saliency (degree 3)*. All specifications control for time fixed effects, controls for day of the week and weekend. Standard errors clustered by product are reported in parentheses. * indicates significance at 10%; ** at 5%; *** at 1%.

Table 8: Overview of Results (effect of saliency on existing products)

	Event day	Event day + 1	Event day + 2	Event day + 3	Event day + 4
Recommended as Substitutes:					
Saliency Shock	0.074***	0.013	0.040**	0.033**	0.010
	(0.018)	(0.017)	(0.017)	(0.016)	(0.018)
Saliency Shock with anticipation control	0.076***	0.017	0.043***	0.034**	0.010
	(0.018)	(0.017)	(0.017)	(0.016)	(0.017)
Saliency Shock with lagged novelty control	0.077***	0.017	0.043**	0.034**	0.010
	(0.018)	(0.017)	(0.017)	(0.016)	(0.017)
Saliency Shock Accounting for Spillovers	0.082***	0.023	0.043**	0.038**	0.009
	(0.018)	(0.017)	(0.017)	(0.016)	(0.018)
Saliency-Price Sensitivity:					
Saliency Shock	0.075***	0.014	0.041**	0.031*	0.008
	(0.018)	(0.017)	(0.017)	(0.016)	(0.018)
Saliency Shock × Price	0.000	0.000	0.000	-0.000	-0.000*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Saliency-Attention Sensitivity:					
Saliency Shock	0.323***	0.133**	0.134**	0.252***	0.277***
	(0.060)	(0.064)	(0.067)	(0.067)	(0.065)
Saliency Shock × Size of Recc. Set	-0.032***	-0.015*	-0.012	-0.028***	-0.034***
	(0.008)	(0.009)	(0.008)	(0.008)	(0.008)
Recommended as Complements:					
Saliency Shock	0.019***	0.001	0.009*	0.004	0.002
	(0.005)	(0.003)	(0.005)	(0.004)	(0.005)

Figure 14: Saliency Shocks: Placebo

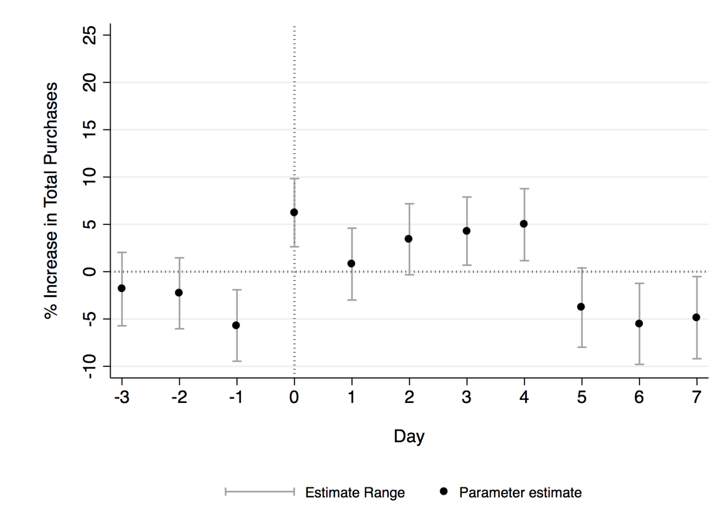

(a) **Saliency America**

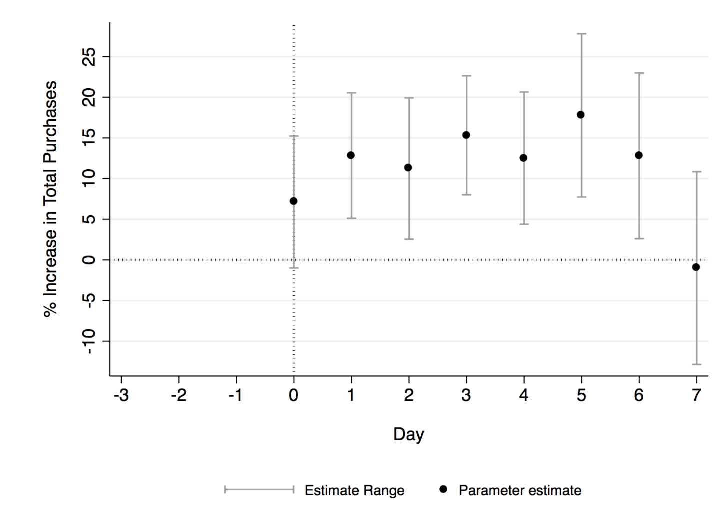

(b) **Saliency America** × **New Product America**

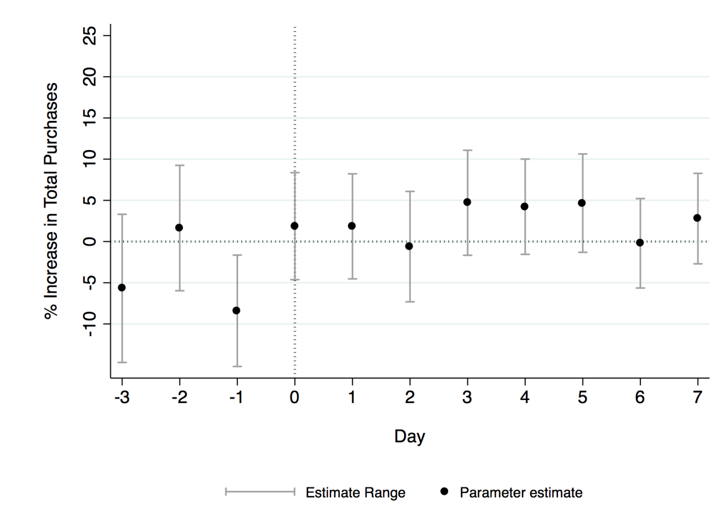

(c) **Saliency Europe** × **No Saliency America**

Table 9: Differences between treatment and control products

	Treatment	Control	Difference	t-statistic
Pre-Event:				
Price (in US$)	1,062.20	1,092.67	30.47	0.5283
US-EUR Diff in Purchase	-0.001	0.012	0.013	0.5283
Post-Event:				
US-EUR Diff in Purchase	-0.050	-0.162	-0.111***	-3.3013
Observations	1296	1773		

∗∗∗ Difference significant at 1%.

Figure 15: Double Diff-in-Diff: Common Trends

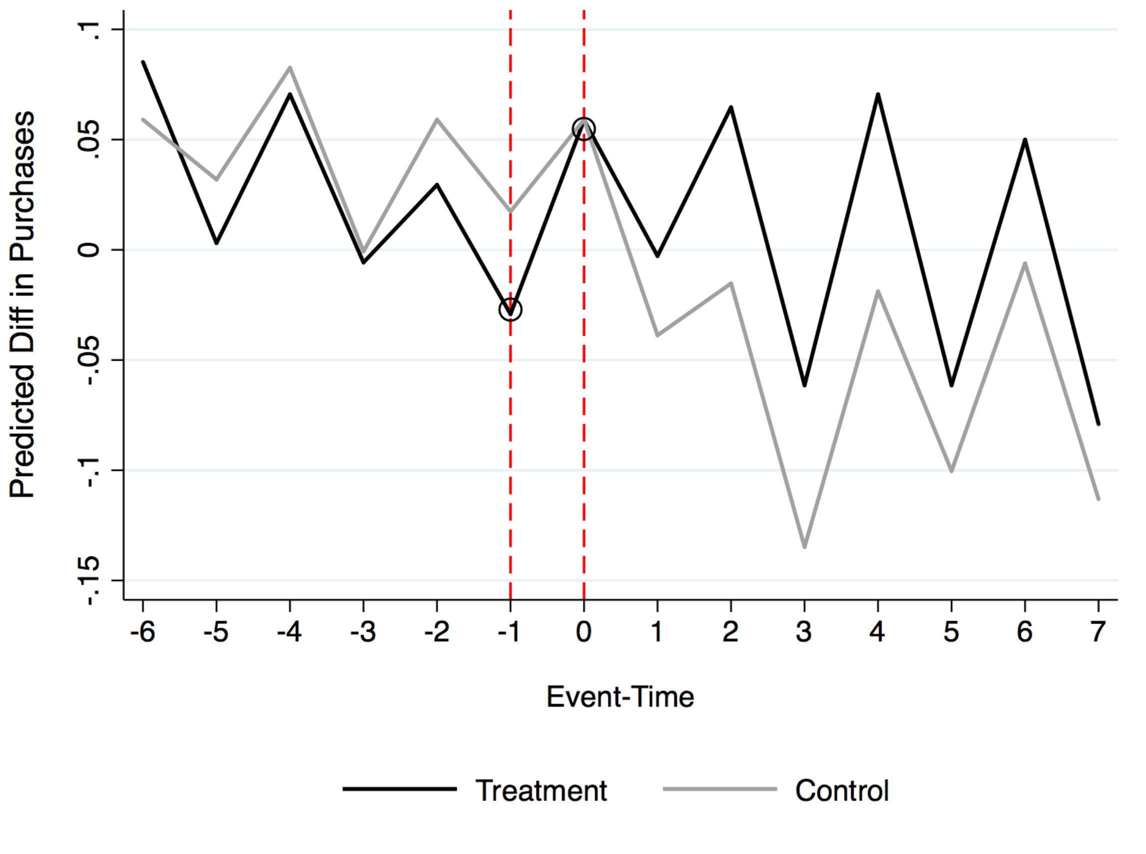

Table 10: Effects of Saliency on Difference in Demand b/w America and EUR

	(1)	(2)	(3)	(4)
Treatment	0.017		0.017	0.000
	(0.029)		(0.029)	(0.029)
Post	0.000		0.155	0.169
	(0.000)		(0.102)	(0.108)
Treatment × Post	0.079**	0.079*	0.244**	
	(0.038)	(0.040)	(0.118)	
Treatment × Post × Size of Set			-0.012*	
			(0.007)	
Post (Day 0)				0.060
				(0.043)
Post (Day 1)				0.075
				(0.058)
Post (Day 2)				0.127***
				(0.049)
Post (Day 3)				0.052
				(0.052)
Controls for product 'age'	Yes	Yes	Yes	Yes
New Product (Block) F.E.	Yes	Yes	Yes	Yes
New Product (Block) × Post F.E.	Yes	Yes	Yes	Yes
Product F.E.	No	Yes	No	Yes
Observations	7161	7161	7161	7161

This table reports results on the effect of the saliency shock on the difference in toal **shopping bag additions** (per day), between America and Europe. The sample consists of a subset of, products that are recommended exclusively in the two regions, America and Europe. For this sample of product the specification estimates a double difference-in-difference equation for the sample's daily transactions over a (-3,+3) event window. *Treatment* is a dummy variable that takes the value 1 if the product was recommended in America but not in Europe. Post a dummy variable indicating the post-event window (0,+3). Our proxy for attention, the *size of the set* is defined as the average size of new product recommendation sets that include the target product. *Age* of the product is the number of days since the product was released in the catalogue for sale. Standard errors clustered by product are reported in parentheses. * indicates significance at 10%; ** at 5%; *** at 1%.

Table 11: Effects of Saliency on Consideration and Choice on Individual Choice (Shopping Bag)

	Travel Bags		Travel Bags		Watches	
	Consideration	Choice	Consideration	Choice	Consideration	Choice
Saliency	**0.1534*****	**0.260*****	**6.733*****	**0.292**	**3.340*****	**-0.532**
	(1.51e-08)	(2.44e-08)	(0.253)	(0.364)	(0.008)	(0.712)
New Product	-2.539***	1.228***	-12.223***	6.422***	-3.284***	2.347***
	(8.96e-08)	(9.92e-08)	(0.470)	(0.280)	(0.797)	(0.393)
Saliency×New Product			-4.687***	-5.166***	0.797***	-1.124*
			(0.253)	(0.293)	(0.008)	(0.637)
Aggregate Marginal Effect of Saliency	0.022		0.030		0.154	
Observations	1080	1080	1080	1080	4430	4430
# Products	12	12	12	12	10	10
# Products with Saliency Shock (at least 1)	10	10	10	10	8	8
# Individuals	90	90	90	90	443	443

This table reports results on the effect of the saliency shock on the probability of purchase, incorporating both consideration and choice. The sample consists of a subset of products (Travel Bags and Watches), purchased in the region America. All specifications control for the product price. *Saliency* is a dummy variable that takes the value 1 if the product was recommended at time t and 5 days following it. The mean marginal effect of saliency is given by (the mean of) the expression: $\left[(1-B_{ijt})\psi_1 \cdot P_{ijt} + \sum_{c_t \in G}\left\{(P_{ijt}|c_t)(1-P_{ijt}|c_t)P_{it}(c_t)\psi_2\right\}\right]$. Standard errors, based on the quadratic approximation to the curvature at the maximum likelihood estimate, are reported in parentheses. * indicates significance at 10%; ** at 5%; *** at 1%.

Table 12: Mean Own and Cross-Elasticities from Change in Saliency

	Elasticities											
Product	1	2	3	4	5	6	7	8	9	10	11	12
1	0.079	-0.021	0.000	-0.005	-0.004	-0.009	-0.007	-0.007	0.000	0.010	-0.006	-0.019
2	0.007	0.068	0.000	-0.005	0.012	-0.007	-0.006	-0.006	0.000	-0.005	-0.005	-0.017
3	-0.007	-0.006	0.000	-0.005	0.005	0.014	-0.005	-0.005	0.000	-0.005	-0.005	-0.017
4	-0.007	-0.021	0.000	0.044	-0.010	0.005	0.015	-0.006	0.000	-0.005	-0.007	-0.019
5	-0.008	-0.020	0.000	-0.005	0.098	-0.009	0.007	0.015	0.000	-0.005	-0.007	-0.019
6	-0.008	-0.020	0.000	-0.005	0.004	0.081	-0.007	0.007	0.000	-0.005	-0.007	-0.019
7	-0.008	-0.021	0.000	-0.005	-0.011	0.005	0.062	-0.007	0.000	-0.005	-0.006	-0.019
8	-0.008	-0.021	0.000	0.010	-0.003	-0.009	0.008	0.062	0.000	-0.005	0.015	-0.018
9	-0.009	-0.021	0.000	-0.004	-0.010	-0.001	-0.007	0.008	0.000	-0.004	0.007	-0.004
10	-0.002	-0.021	0.000	-0.005	-0.010	-0.007	-0.006	-0.007	0.000	0.044	-0.007	-0.005
11	0.013	-0.020	0.000	-0.005	-0.011	-0.009	-0.006	-0.006	0.000	-0.005	0.062	-0.019
12	0.007	0.002	0.000	-0.005	-0.010	-0.008	-0.006	-0.005	0.000	-0.005	0.009	0.048
Average Elasticity												
Spillovers	-0.006	-0.003	-0.003	-0.005	-0.004	-0.005	-0.007	-0.003	-0.004	-0.006	-0.006	-0.002
Own+Spillovers	0.004	-0.010	0.000	0.001	0.004	0.004	0.004	0.004	0.000	0.001	0.004	-0.010

Cell entries j, k, where j indexes row and k column, give the percent change in sales share of brand j with a one-percent change in the saliency of product kj. Each entry represents the mean of the elasticities from the 90 consumers. The last two rows report the average percent change in sales share of all products due to a one-percent change in the saliency of product j. The penultimate row ('Spillovers') omits self-elasticities in computing this average.

Figure 16: Counterfactuals - Limited Attention vs. Full Attention Difference in Sales Share

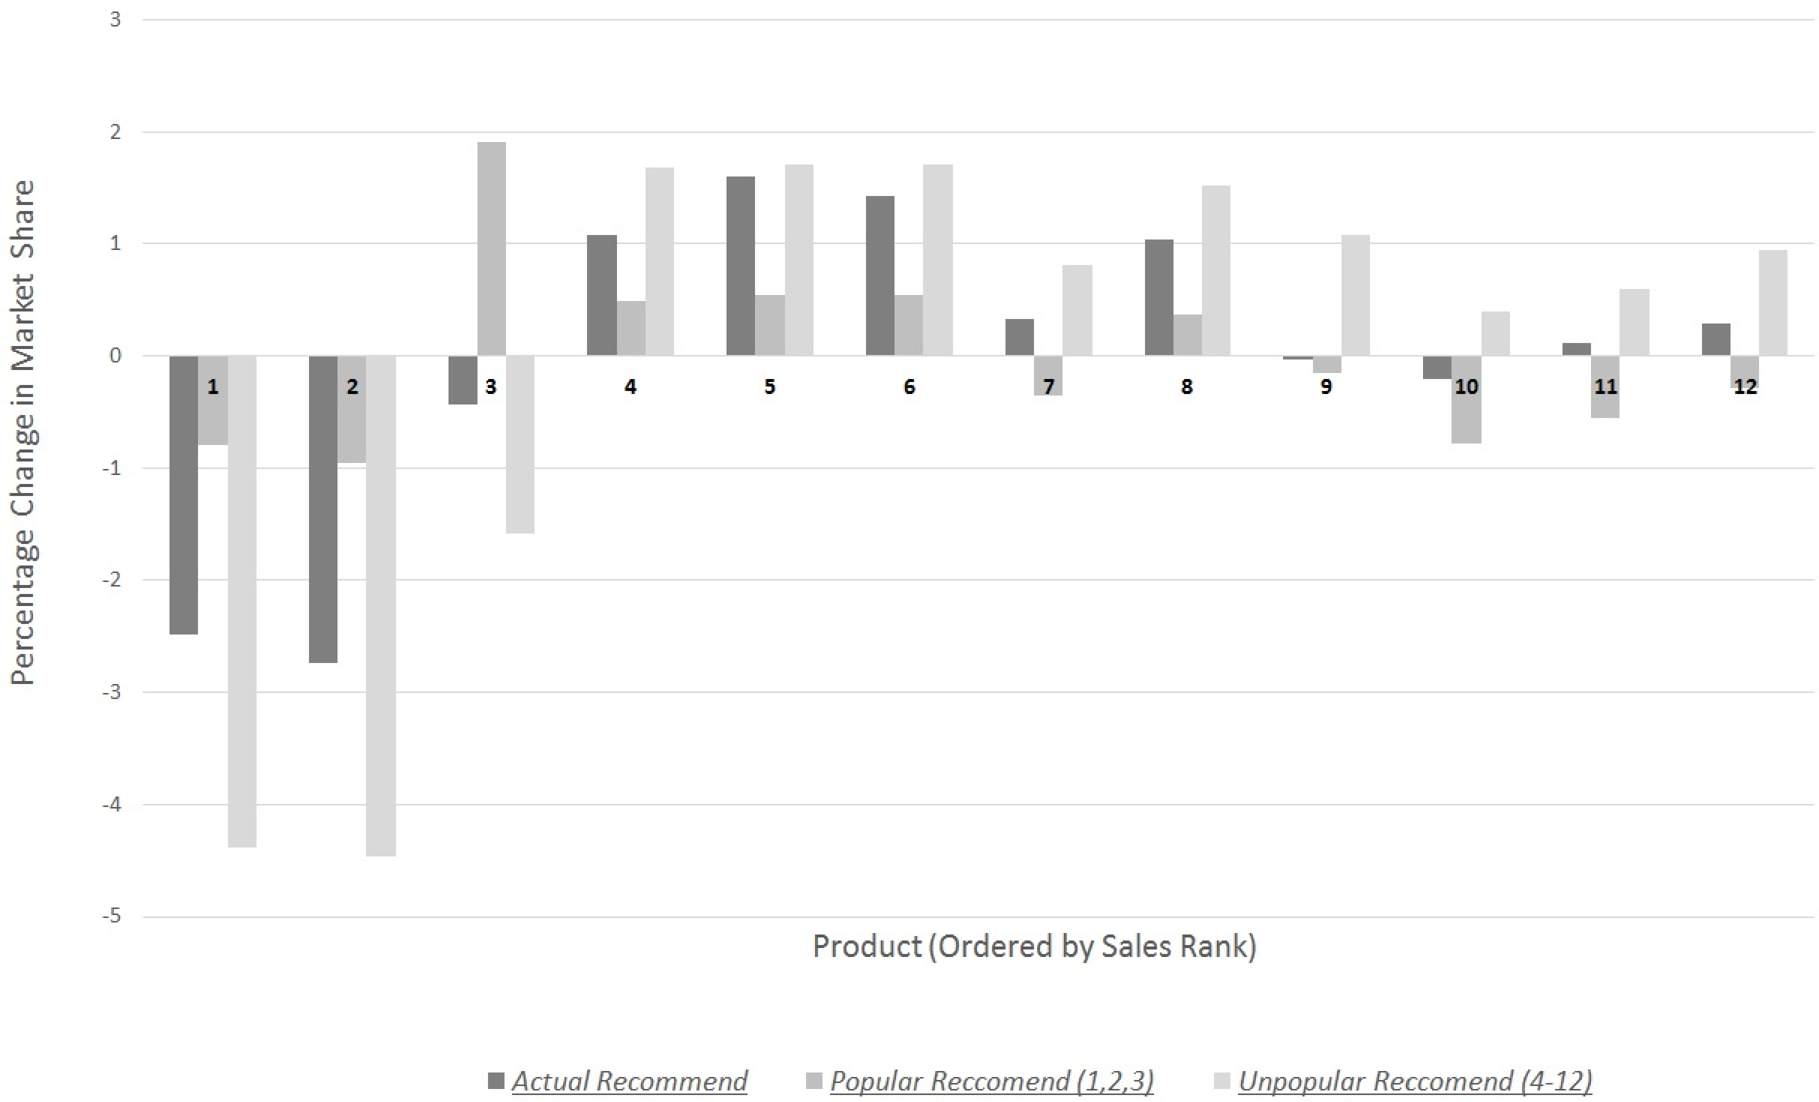

The figure shows results from counterfactual estimations allowing for the following scenarios: (1) consumers have full, instead of limited, attention and (2) only certain types of products receive recommendations. The x-axis orders products by their sales rank (1 being most popular and 12 being least popular). Products 1, 2 and 3 are categorized as 'popular products' (observed sales share > 10%); products 4 to 12 are categorized as 'unpopular products.' Each bar represents the percentage difference in sales share between when consumers have limited attention and when they have full attention. Dark gray bars show results under the existing recommendation system. Gray bars show the difference between limited and full attention when only popular products are (always) recommended (products 1, 2 and 3). Light gray bars show the difference between limited and full attention when only unpopular products are (always) recommended (products 4 to 12).