Publication information

Title	When Will Workers Follow an Algorithm?: A Field Experiment with a Retail Business
Author(s)	Kawaguchi, Kohei
Source	Nil
Version	Nil
DOI	Nil
Publisher	Nil

Copyright information

© The Author

Notice

This version is available at HKUST Institutional Repository via

<http://hdl.handle.net/1783.1/100684>

If it is the author's pre-published version, changes introduced as a result of publishing processes such as copy-editing and formatting may not be reflected in this document.

For a definitive version of this work, please refer to the published version.

<http://repository.ust.hk/ir/>

When Will Workers Follow an Algorithm?: A Field Experiment with a Retail Business *

Kohei Kawaguchi[†]

September 1, 2019

Abstract

This paper develops a new algorithm for increasing the revenue in a dynamic product assortment problem. Then, it identifies the challenges faced by managers in practice and discusses the conditions under which workers follow the algorithm. To do so, I conducted a field experiment with a beverage vending machine business. The experiment shows that, on average, workers are reluctant to follow the algorithmic advice; however, the workers are more willing to conform once their forecasts are integrated into the algorithm. Analyses using non-experimental variations highlight the importance of taking worker and context heterogeneity into account to maximize the benefit from adopting a new algorithm. Higher worker's regret, sales volatility, and fewer delegations increase the conformity, while they mitigate the effects of integration. Workers avoid high-traffic vending machines and focus on machines with high sales volatility when adopting the algorithm. The effects on the sales are largely similar to the effects on product assortments. The results emphasize the gap between nominal and actual performance of an algorithm and several practical issues to be resolved.

Keywords: Automation; Experimentation; Learning; Product Assortment Problem; Multi-Armed Bandit Problem; Algorithm Aversion; Conflicts of Interest; Productivity; Retail Business; Beverage Vending Machines

JEL Codes: L20; L81; M12; M15; M20.

*This work was supported by JSPS KAKENHI Grant Number 16K17111 and by the Japan Center for Economic Research grant. I thank Kosuke Uetake, Yasutora Watanabe, Andrew Ching, and seminar participants at East China Normal University, Musashi University, Chinese University of Hong Kong, Xiamen University, Nazarbayev University, Hong Kong Baptist University, The University of Hong Kong, Rice University, Southwestern University of Finance and Economics, and Asia-Pacific Industrial Organization Conference for their valuable comments. I thank Yuta Higashino, Masato Honma, Takayuki Konno, Shunsuke Ijima, Takeshi Kato, Takayuki Ishidoya, Toshinari Sasagawa, Takatoki Izumi, Yukinori Kurahashi, Tomoki Mifune, and many others at JR East Water Business as well as its operators for their support and insightful conversations throughout the project. I thank seminar and conference participants for their constructive comments. I thank Enago (www.enago.jp) for the English language review.

[†]Department of Economics, School of Business and Management, Hong Kong University of Science and Technology. e-mail: kkawaguchi@ust.hk

1 Introduction

Algorithms are taking over decision-making abilities of humans in every business domain. However, achieving full automation of a business activity is rather an exception, which includes the Internet companies, in which programs also implement algorithmic decisions, and manufacturers, in which robots are able to work in a well-controlled environment. In several other situations, such as retail industries operating in a real-world shopping site, decisions regarding pricing, product assortment, and inventory control, must be physically implemented in an environment that is extremely complicated to automate. With partial automation, a manager at best can give workers algorithmic advice, expecting that the workers comply with the advice.

However, the existing literature already provides some evidence that a strategy of simply ordering workers to follow an algorithm may not work. A phenomenon in which people often obey inferior human decisions, even if they understand that algorithmic decisions outperform them, is widely observed. This is known as *algorithm aversion* (Dietvorst et al., 2015). Algorithm aversion may prevent workers from following algorithms. Similarly, Hoffman et al. (2018) showed, in the context of hiring, that managers often overruled machine-based recommendations, which resulted in worse outcomes. The implementation cost also leads to conflicts of interest between a company and its workers. Such agency costs may also prevent workers from conforming to the algorithm, especially if the implementation of a task under the algorithmic advice increases the intensity of work. Although not an algorithm, for example, Atkin et al. (2017) identified conflicts of interest as organizational barriers to technology adoption for the production of soccer balls in Pakistan. However, the literature also proposes a way of increasing the conformity to algorithmic advice. In a laboratory experiment, Dietvorst et al. (2016) demonstrated that people were more willing to follow algorithms if only a small amount of control over the decision was allowed.

In this paper, I develop a new algorithm for increasing the revenue in a dynamic product assortment problem. Beyond the usual theoretical regret analysis and simulation analysis, I attempt to identify the challenges encountered by managers when they put the algorithm into practice. Namely, I study the conditions under which workers are more or less likely to follow the algorithmic advice and the automation can be successfully implemented. In particular, by means of a field experiment, I investigate whether integrating a worker’s opinion into the algorithm can increase the worker’s conformity to the algorithm, as in Dietvorst et al. (2016). Then, I analyze the heterogeneity of workers’ responses to the algorithmic advice by utilizing non-experimental variations in the worker and business area characteristics such as the performance of the worker in the previous season, degree of delegation to the worker,

experience of the worker, sales volume, and volatility in the area.

My investigation takes the form of a field experiment in a beverage vending machine company operating in the Tokyo metropolitan area. Among many tasks performed by the workers in this company, I focus on the task of experimentation with product assortments facing demand uncertainty. In this company, a worker, who is hired by a third-party subcontractor, is assigned to a narrowly-defined area that comprises a few dozen vending machines. The worker is responsible for maintaining and stocking the machine, selecting some of the product assortment in the vending machines slots. The workers are incentivized to experiment with products and adapt the assortments to local demand.

I focus on the product assortment problem because it is one of the most important and computationally demanding decision variables in retail businesses as much as pricing. Following recent developments in the literature of multi-armed bandit problems (for example, see Bubeck and Cesa-Bianchi (2012) for surveys of the existing results), several companies will start to introduce this type of algorithms into their business. Therefore, it is important to study the gap between nominal and actual performance of an algorithm and the practical issues to be resolved.

First, I develop an algorithm to automate the experimentation in the real business setting of the company. It starts from a prior belief of the demand parameters of products in a particular area during a business season. The algorithm chooses the product assortments each week for each vending machine in an area. At the end of the week, it observes the weekly data and updates the belief, continuing the process until the end of the season. It balances the exploitation and exploration according to a certain rule. Because this task lacks a readily available algorithm and it is not possible to derive an optimal policy using dynamic programming, I arrange a Thompson Sampling (TS) algorithm. This algorithm is known to work well under settings like the one in this experiment. Recently, Abeille and Lazaric (2017) analytically showed the regret boundaries on a wide class of TS algorithms. The algorithm developed for this study can be considered as a special case of their TS algorithm for generalized linear context bandit problems. According to their results, the algorithm below has a regret bound of rate $O(\sqrt{T})$. With a simulation, I demonstrate that this algorithm substantially increases the total revenue, assuming that the demand function is correctly specified.

Then, I conduct a field experiment, in which I introduce the previous algorithm into the business of the company for a season. In the experiment, I determine i) if the workers follow the algorithmic advice, and ii) if they are more willing to follow the algorithm when it integrates workers' own forecasts. I randomly divide the areas into three groups. In one treatment group, I provide advice on the basis of the previous algorithm not using workers'

forecasts. In another treatment group, I give advice based on the algorithm integrating the workers' forecasts for the demand of the new products into the prior belief of the algorithm, and I inform the workers whether their algorithm integrates their forecasts. Because the output of the algorithm is nearly the same with and without the integration of workers' forecasts, any difference in the outcome is caused by the mere perception of the workers that the algorithm integrates their forecasts.

I find that the product assortments in the first treatment group are no more correlated with the output of the algorithm than those in the control group. In the control group, the algorithm outputs the optimal product assortments, but they are not shared with the workers. This indicates that workers, on average, are reluctant to follow the advice based on an algorithm. The follow-up survey identifies some of the reasons: some of the workers perceived it increases their burden of work, did not think that the demand prediction was reliable, and felt inconsistent with the other instructions given by the company.

In contrast, the product assortments of the second treatment group are correlated more with the algorithm's output than those of the first treatment group. Thus, integrating workers into the design of the algorithms may resolve a part of the problem. My findings are consistent with the results from existing laboratory experiments (Dietvorst et al., 2015, 2016). However, the magnitude of the impacts on the average product assortment behaviors is not substantial. I find that the effects of algorithmic advice on revenue are not statistically significant either in the first or second treatment group.

However, different types of workers may respond differently to the introduction of an algorithm. If this is true, we may still be able to improve the performance by carefully choosing subsets of workers who are better able to comply with an algorithm or by changing the method of communication with the workers. To explore these possibilities and to better understand when and why workers follow an algorithm, I next study the heterogeneous responses of the workers using non-experimental variations in worker and business area characteristics.

The results show several interesting patterns for area/worker-level heterogeneity. I find that a worker is more likely to follow the algorithmic advice if the worker has higher regret compared to the TS algorithm in the last business period, faces lower sales volatility in his area, or has a lower degree of delegation. The work experience of a worker is not significantly correlated after controlling for the above factors. However, integrating a worker's opinion makes the worker more likely to follow the algorithmic advice if the worker has lower regret compared to the TS algorithm, faces higher sales volatility, has more work experience, or has higher degree of delegation.

I also find that workers carefully select which vending machines to use for their experi-

ments with the algorithmic advice. First, a worker is more reluctant to follow the algorithm in machines with high sales volume. A worker may fear the risk of experimenting with a new assortment policy in high traffic machines. Second, a worker is more willing to follow the algorithm in machines with high sales volatility. A worker may find it difficult to adapt to the demand alone in such a machine and would rely on the algorithm. Third, the integration is more effective in machines with high sales volume and low sales volatility. For these types of machines, workers are confident in their decisions and need more justification for following the algorithm.

Workers' actual product assortments and TS assortments are different in several ways. In particular, I find that workers tend to allocate higher slot shares than the algorithm to new products, private brands, hot beverages, large products, cheap items, and Coca Cola products. When algorithmic advice is given, the differences are statistically significant and become smaller. Thus, the misallocation of vending slots compared to the TS assortments across product types are adjusted under the algorithmic advice, even though the overall correlation was not significantly different from the control group because of the size of the effects. I find that integrating workers' opinions further adjusts the assortments, especially for the share of hot and cold beverages. The difference in the impact on sales between the algorithmic advice with and without integration seems to be driven by the allocation of hot and cold beverage slots. This makes sense because the timing of introducing hot beverages when the season changes from summer to winter has been recognized as one of the decisions that matters most in the sales, among managers of the company.

The effects on the sales are significant in a few cases, and the coefficients signs are largely in line with the signs of the corresponding coefficients in the effects on product assortments. Thus, if the manager cannot directly control workers behaviors, it is important to understand the above heterogeneous responses of workers to maximize the potential benefit from adopting an algorithm. The results highlight that there is a distance between nominal and actual performance of an algorithm once it is practically applied. At the same time, it indicates that there is a systematic pattern in the way workers more or less follow the algorithmic advice. The conditions highlighted in this paper would help managers design the way workers and algorithms are integrated in their business.

The remainder of the paper is organized as follows. In the rest of this section, I review related literature. In Section 2, I describe the industry and institution with summary statistics of the data. Section 3 estimates the beverage demand function. Section 4 describes the algorithm for the current product assortment problem and demonstrates the effectiveness using simulations. Section 5 reports the results of the field experiment. The last section concludes with implications and limitations of this paper.

1.1 Related Literature

In the field of forecasting, it has been widely recognized that people often put more weights on humans’ forecasts than algorithms’ forecasts. For example, in a laboratory experiment, Önköl et al. (2009) found that people adjust their forecasts more when the advice is perceived to come from a human expert than from a statistical forecasting method. Dietvorst et al. (2015) found that people more quickly lose trust on the algorithmic than human forecasters when they make the same mistake. Dietvorst et al. (2016) found that the algorithm aversion could be eased by ceding a small amount of control. Prah and Van Swol (2017) theorise the phenomenon in the psychological framework. Fehr et al. (2013), Bartling et al. (2014), and Owens et al. (2014) theorize the value of authority and control in the principal-agent framework.

To the best of my knowledge, my paper is the first to test the algorithm aversion and the effectiveness of integrating workers’ opinions in the actual business setting outside laboratories. Dzindolet et al. (2002) showed that there is a discrepancy in the self-report and performance data regarding the use of automated systems. There can be a difference between laboratory and field experiments, because the use of algorithm actually affects their rewards. The paper also differs from the literature in studying the introduction of multi-armed bandit algorithm, which solves much more computationally demanding task than forecasting.

For this research, I build an algorithm to automate the experimentation of product assortment decisions in the beverage vending machine business, on the basis of the ideas developed for the "multi-armed bandit" problem literature. I apply TS (Thompson, 1933) to the vending machine-level product assortment decisions across several weeks. Several studies, like Scott (2010), Chapelle and Li (2011), and May et al. (2012), demonstrated the empirical efficacy of TS vis simulations. Agrawal and Goyal (2012) showed that the TS algorithm achieved logarithmic expected regret for the stochastic multi-armed bandit problem. Caro and Gallien (2007) and Rusmevichientong et al. (2010) studied the dynamic product assortment problem and proposed heuristic policies to deal with it. However, their policies were not directly applicable because they do not consider a multiple-site situation. Additionally, Caro and Gallien (2007)’s policy exploited a specific form of demand that had no substitution across products. Thus, it was not empirically appropriate.

Several studies such as those by Israel (2005), Erdem et al. (2005), Narayanan et al. (2005), Chernew et al. (2008), Erdem et al. (2008), Narayanan and Manchanda (2009), Ching (2010), Anand and Shachar (2011), Szymanowski and Gijbrecchts (2012), and Ching et al. (2013), used consumer choice data to estimate consumer learning models and examine its implications for marketing and policy interventions. A similar approach was employed to study the efficacy of physician’s learning by Coscelli and Shum (2004), Crawford and Shum

(2005), and Ching and Ishihara (2010). There are several papers that theoretically studied the experimentation by firms, including Rob (1991), Mirman et al. (1993), Creane (1994), Raman and Chatterjee (1995), and Hitsch (2006) in a monopoly market, and Aghion et al. (1993) and Harrington (1995) in a duopoly market, and Harpaz et al. (1982), Bertocchi and Spagat (1998), and Bergemann and Valimaki (2000) in a more competitive market. They mostly studied price experimentation, whereas the current paper is about experimentation with product assortment decisions. Additionally, my paper is empirical. Bolton and Harris (1999), Keller et al. (2005), and Keller and Rady (2010) studied the implication of information spillovers during strategic experimentation in an abstract setting. This paper does not estimate the experimentation and learning process, but it studies the responses of workers when an algorithm for these tasks is introduced.

2 Industry and Institutional Background

2.1 Industry Background

Vending machines are important retail channels in Japan particularly for the beverage business. According to the Japan Vending Machine Association,¹ there were 2.6 million installed beverage vending machines at the end of 2013. That year, vending machine beverage sales totaled 2.3 trillion yen. This is approximately one third of the total annual beverage sales in Japan and is equivalent to supermarket sales.

A unique feature of the industry is that retail prices are constant over time and locations except for unusual locations, like cinemas or mountaintops, and only differ across product packages. The nominal price changed only when the consumption tax was raised from 3% to 5% in 1997 and from 5% to 8% in 2014. Between 1997 and 2014, the price of a 350 mL can was 120 yen, and the price of a 500 mL bottle was 150 yen across the industry. Price variations across products are only exceptional within package types.²

2.2 Institutional Background

Japan Railway East (JRE) is the largest railway company in Japan. Originally part of the national railway, it was privatized in 1987. The company operates in the eastern part of mainland Japan including the Tokyo metropolitan area. In addition to the transportation

¹<http://www.jvma.or.jp/index.html>.

²The reason for this price rigidity and homogeneity across locations and products is outside the scope of the paper. Possible explanations for time-in variance include the low inflation in the 2000s in Japan and menu costs. Homogeneity across locations may be attributed to retail price maintenance motivated by fire-sale fears among competitive retailers (Deneckere et al., 1996, 1997).

business, the company operates a retail business inside stations to exploit the substantial number of passengers who regularly use the transportation service.³ The beverage vending machine business is a branch of that retail business unit. The JRE railway lines and stations in the Tokyo metropolitan area are shown in Figure 1.

There are multiple models of vending machines installed in JRE stations with between 30 and 42 slots. The most common machine model has 36 slots. Figure 2 shows a picture of a typical vending machine. Popular brands can occupy multiple slots, but in most cases one brand occupies only one slot. A package of available products is displayed on the upper half of the front panel. Consumers choose the item they want to buy by pushing a button under the package. The consumer then inserts cash or touches the sensor with a JRE electric commuter card called SUICA in the middle of the panel. Then, the item drops into the box at the bottom of the machine. Because almost all passengers in the Tokyo area have this commuter card, and it is used for a large percentage of transactions. The system and machinery are almost identical to vending machines of other companies outside stations. For each slot displayed on the panel, there is one box for stock inside the machine, which, on average, stores 30 cans or bottles. The remaining stock is not visible to consumers. If a brand stocks out, then letters will appear at the bottom of the mock package, indicating its status.

JRE has a subsidiary firm managing the company’s beverage business. This is the principal company in my analysis. The company owns all the vending machines in JRE stations. The firm outsources maintenance to third-party subcontractors. In addition to maintenance, the principal company delegates some of the authority related to product assortment decisions for vending machines. On average, product selection for up to 70% of the slots in the vending machines is directly determined by the principal company. However, assortment for the remaining slots is the responsibility of the subcontracted workers. I call the first slots the *centralized slots* of a vending machine, and the latter the *delegated slots*. The principal company expects and encourages the subcontractors to use these slots for exploring and adapting to local demand. Because the product assortment policy is particularly important for my analysis, I explain JRE’s assortment policy in further detail in the next section.

Subcontractor operations are divided into several local areas. A local area, on average, consists of three stations closely located, mostly along the same line with similar demographic characteristics. A local area is, in principle, operated by a single staff member or a small fixed number of staff during a season. This is the basic decision-making unit in the business. Thus, I focus on their decisions with this study. Figure 3 shows a picture of a worker refilling

³The average number of daily passengers in a station in Tokyo ranges from tens to hundreds of thousands: <http://www.breast.co.jp/passenger/>.

products at a vending machine at one of the stations.

The fees paid to the subcontractors are linear in the sales from the allocated vending machines. Therefore, workers are concerned with top-line sales, rather than profits. A representative of the principal company explained that this is because there is little variation in margins except for private brands, whose assortment is completely determined by the principal. The fee rates differ across subcontractors by a few percentage points. Unfortunately, this information is confidential and cannot be used for the analysis. The industry average of the fee rate is 20% of the total sales per an industry report (Morita et al., 2012). The subcontractors bear all operational costs, which can potentially create the conflict of interests. At the end of the season, if stocks remain, JRE may buy back the stock but only if they are JRE private brands. The subcontractors are responsible for the stock of non-private brands.

Workers regularly visit the vending machines a few times a day, down to a few times a week, to refill stock, depending on the market size and throughput. Thus, an out of stock vending machine is rare.⁴ Because workers regularly visit vending machines regardless of product changes, marginal costs of changing products atop driving, visiting and refilling are not high.

Products are classified into 12 categories: Green tea, English tea, Other tea, Water, Tonic water, Canned coffee, Bottled coffee, Carbonated water, Fruits, Healthy drink, Specialty, and Other. The principal company defines the categories and classify products. The company has used the same classification during the entire data period. Other tea includes tea products such as Oolong tea, barley tea, and herb tea. Tonic water includes isotonic drink to supply energy. Carbonated water includes soda drinks with different kinds of flavor. Fruits includes fruits juice such as orange, apple, and peach juice. Healthy drink is called by this name by the company but should be considered as energy drink. Specialty drinks are beverages whose effects on the body health is certified by the government. Other includes experimental and special drinks such as corn soup, nata de coco, and jelly.

2.3 Product Assortment Policy

The business plan is determined for each season. April to September is the Spring/Summer, and remaining months are the Fall/Winter. By the beginning of April or October, the principal company selects a list of products consisting of approximately 200 brands. Products to be inserted into the vending machines are chosen from this list. No other products can be inserted. Some products are launched during a season. Until then, the product cannot

⁴Principal company has a system to monitor the time of stock-out for each machine. According to this database, the time of stock-out is less than 2% per day even at the worst one percentile of vending machines. Because of this I abstract away from stock-out events in the empirical model.

be selected. The release dates are announced in advance.

Furthermore, the principal company determines the central product assortment policy. The principal company classifies the vending machines into 31 groups according to the machine model and its sales volume. During my data collection period, this classification was revised twice between the Fall/Winter 2014 and Spring/Summer 2015, and between the Spring/Summer 2016 and the Fall/Winter 2016. The principal company chooses a subset of products from the list to be put on the vending machines belonging to each group. Thus, the central policy is the same across vending machines in a group and differs only between groups. Required products are chosen on the basis of the relative sales volume and growth within a product category in the last seasons. It also changes significantly between Spring/Summer and Fall/Winter and changes slightly across October, November, and December-March within the Fall/Winter period, and across April-June and July-September within the Spring/Summer period.

The centralized product assortment policy determines the products in the centralized slots. The history of product assortment is monitored by the principal company. The principal company annually evaluates the consistency of the worker’s assortment decisions with the centralized plan as an independent item. This result does not directly affect the reward, but is publicly shared. The data confirm that the workers closely follow the centralized plan without significant delay. For the delegated slots, workers are free to select products. The principal company expects that local workers will utilize them to explore and adapt to local demands.

2.4 Summary Statistics

This section summarizes the data. Table 1 displays the number of unique values of products, vending machines, stations, areas, and subcontractors in the sample. Here, I distinguish hot and chilled beverages of the same brand to calculate the unique value of the products. On average, there are 8.89 vending machines per station, and 26.42 per area. Table 2 summarizes the machine-product-week-level sales units and values. The average unit sales is 17.924 and the value is JPY2455.320. There are a few vending machines that have a huge market. They are installed in terminal stations such as Tokyo, Shinjyuku, and Shinagawa. Hundreds of thousands of passengers use these stations every day. Table 3 summarizes the weekly number of newly added products, compared with the prior week’s assortments for the machine-week level during a season. It only counts products put on delegated slots. This shows that workers are frequently changing products at the vending machine level. In the first 10 weeks, the number of newly added products are kept above 2.0. This number is non-negligible because

the number of delegated slots is approximately 10 per machine. The number increases during the first week; then it drops to around 0.5. This is the typical pattern of experimentation.

Table A1 in the online appendix⁵ shows the summary statistics of product assortments per machine. The most common vending machine has 36 slots and at the median carries 33 unique products. This means that typically a machine of this type assigns two slots to two products and one slot to other 31 products or assigns three slots to one product. Thus, a product rarely occupies more than one slot. A vending machine covers 11 categories even at the 1st quartile. Hot products occupy on average 23.7% of the slots. Figure A1 displays the weekly trend in the share of hot product slots. The share is around zero between 20th and 40th week and quickly expands after 40th week. Over years, the timing of reducing the share of hot products has been advanced by a few weeks. Table A2 shows the summary statistics of product sales per machine and this shows that the share of hot product sales is on average 27.3%. This may suggest that the company better slightly expands the share of hot product slots. The tables also shows that the share of new product is 23.7% for slots and 20.2% for sales. This indicates that the company is exploring new products at the cost of current revenue. The share of private brand products is almost negligible for both slots and sales. More recently, the company has gradually increased the share of private brand products, but it is another story. The tables also show the share of slots assigned to products that are not required by the principal company. Here, a slot is said to be a delegated slot if the product assigned to the slot is not a requirement of the season. Because at the beginning of a season, products required in the last season still stay in vending machines, the share of delegated slots calculated in this way tend to be higher than the nominal share of 30%. Products required by the principal company are uniformly and monotonically replaced with products required in the season over time. Throughout years, on average, 49.8% of slots are filled with products not required in the season. The share of sales from delegated slots is 46.7%, which is lower than the slot share. Table A3 and A4 show the slot and sales share of each product category per machine. It shows that canned coffee products occupies the largest 16.4% of the slots and carbonated water products follows with 14.0%. Specialty products occupies the lowest 1.7% shares. Other products occupy 4-13% of the slots. Thus, the product assortments are on average diverse and not concentrated on a few product categories. The sales shares of each category are largely the same with the slot shares.

⁵Hereafter, the tables and figures in the online appenedix are prefixed with A.

3 Beverage Demand

3.1 Specification

In this section, I estimate the demand function and obtain *ex post* belief about the demand parameters. A pairing of local area and season is indexed by $i \in \mathcal{I}$. In each local area and season, there are $\mathcal{T}$ weeks, $\mathcal{K}$ vending machines, and $\mathcal{J}$ available products. The products are partitioned into $\mathcal{G}$ groups, and the set of products belonging to group $g \in \mathcal{G}$ is denoted by $\mathcal{J}_g$. A group refers to a product category I explained in the previous section. These sets in general depend on i and t . However, I suppress the dependency for notational simplicity.

I assume that demand at a vending machine follows a nested multinomial logit choice model. Thus, the choice specification is made for a practical reason. In the following analysis, I solve the optimal product assortments given a demand function for many occur. When the demand function is a nested-logit form, I can utilize the algorithm of Gallego and Topaloglu (2014) to solve a static product assortment problem with a group-by-group slot restriction. One may think that the demand function should be a mixed-logit form, but Rusmevichientong et al. (2014) showed that consumer-level unobserved heterogeneity made the problem NP-hard. Another possibility is that not only sellers but also buyers are learning the quality of the new products. I am not aware of any theory to pin down the optimal behavior of sellers in such a situation or an algorithm that efficiently solves the problem. I show that the current nested-logit demand function fits the data very well. However, I cannot completely rule out the misspecification.

The product set $\mathcal{J}$ is partitioned into product categories $\mathcal{J}_g, g \in \mathcal{G}$. The utility of consuming product j for a consumer at week t is $v_{ijt} + \nu_{ijt}$, where v_{ijt} is the mean indirect utility, and ν_{ijt} is a preference shock. Let v_{i0t} be the mean indirect utility of not purchasing, normalized to 0. It is assumed that the preference shock has cumulative distribution:

$$\exp \left[- \sum_{g \in \mathcal{G}} \left(\sum_{j \in \mathcal{J}_g} \exp(-\nu_{ijt}/\gamma_i) \right)^{\gamma_i} \right],$$

where γ_i is a dissimilarity parameter. Under this assumption, the preference shocks are uncorrelated across product categories, but can be correlated within a group. The correlation within a group is approximately $1 - \gamma_i$. A higher value of γ_i means greater independence and less correlation. A sufficient condition for the model to be consistent with random utility maximization is imposing $\gamma_i \in (0, 1]$. At $\gamma_i \rightarrow 1$, the model converges to a multinomial logit choice.

In this specification, the choice probability of product j belonging to group g in week t

at vending machine k , s_{ijkt} , is:

$$s_{ijkt} = \frac{\exp(v_{ijt}/\gamma_i)d_{ijkt}}{\sum_{l \in J_g} \exp(v_{ilt}/\gamma_i)d_{ilkt}} \frac{\left(\sum_{l \in J_g} \exp(v_{ilt}/\gamma_i)d_{ilkt} \right)^{\gamma_i}}{1 + \sum_{m \in \mathcal{G}} \left(\sum_{l \in J_m} \exp(v_{ilt}/\gamma_i)d_{ilkt} \right)^{\gamma_i}},$$

where d_{ijkt} denotes product assortment decisions taking the value of 1 if product j is selected for vending machine k in week t , and takes the value of 0 otherwise. I let $s_{ij|gkt}$ be the choice share of product j among available products belonging to group g in week t at vending machine k . The model can be transformed into a linear regression model:

$$\ln(s_{ijkt}) - \ln(s_{i0kt}) = v_{ijt} + (1 - \gamma_i) \ln(s_{ij|gkt}).$$

It is assumed that the indirect utility of a product has the form of:

$$v_{ijt} = \xi_{ij} + \beta_{ig} \times \text{Temperature}_{it},$$

where ξ_{ij} is a product-specific unobserved intercept and β_{ig} is the group-specific sensitivity to temperature. Because the price of a product is fixed across location and time, the income effect term is absorbed into ξ_{ij} . I call ξ_{ij} s, β_{ig} s, and γ_i , *demand parameters* of a local area in a season.

I do not impose any restriction to the demand parameters across areas and seasons. Hence, I estimate the parameters separately for each local area and season. To do so, I consider the following linear normal model:

$$\ln(s_{ijkt}/s_{i0kt}) = \xi_{ij} + \beta_{ig} \times \text{Temperature}_{it} + (1 - \gamma_i) \ln(s_{ij|gkt}) + \sigma_i \epsilon_{ijkt},$$

with an i.i.d. standard normal random variable ϵ_{ijkt} , generating posterior samples of ξ_{ij} s, β_{ig} s, and γ_i for each local area and season with a multivariate normal prior for ξ_{ij} s, β_{ig} s, and γ_i , and an inverse Gamma prior for the error variance. The prior mean of ξ_{ij} s, β_{ig} s, and $1 - \gamma_i$ are set to 0, and the standard deviation is set to an arbitrary large number, 1000. The shape and slope parameters of the inverse Gamma prior are 0.005. I obtain a 10,000 sample and burn-in the first 5,000. I obtain the sample for each local area and season and compute the *ex post* belief mean.⁶

⁶If a product is not observed in a local area in a season, as a convention, I use the latest *ex post* belief in the local area. If a product has not been observed in the local area, I use the average of the *ex post* belief mean of the same product group in the local area in the season.

To calculate the choice share of a product at a vending machine in a week, s_{ijkt} , I use the sales count data. The area size of a vending machine is assumed to be 100 times its average weekly total sales. The choice of multiplier only affects the location and scale of the demand parameter. To calculate the choice share of a product among available products belonging to the same product group, I use the actual product assortment data.

3.2 Estimation Results

The model fits the data well. The left panel of Figure 4 displays the actual versus the fitted sales share in the log difference from the outside share under the *ex post* belief of the demand parameters. Because the sample size is huge, I randomly sampled 10,000 observations to display the plots. The R^2 of the model is 0.658. Because I use the demand model to forecast future sales, I also check the one-period-ahead forecasting accuracy. I estimate the demand model using data up to period t and forecast the sales in period $t + 1$ for each period. The right panel of Figure 4 displays the actual versus the one-period-ahead forecast. The one-period-ahead R^2 is 0.453.

Figure A3 plots residuals against the lagged residuals from the demand function estimation. The left panel shows residuals against 1-week lagged residuals and the right panel shows against 4-week lagged residuals. It shows that there is a significant serial correlation in the residuals, whereas they become mild on a monthly basis. Thus, we can potentially improve the prediction performance of the model by including the current residual estimates. However, I discard this information in the following algorithm I use in the field experiment, because it is not straightforward to incorporate the uncertainty in the residual estimates into the algorithm.

The summary statistics of the *ex post* belief across areas and seasons are shown in Tables A5, A6, and A7. The dissimilarity parameter γ_i is on average 0.751 with a standard deviation 0.089. Because it is below 1, there is some correlation in the preference shock within a product category. However, correlation is mild because the parameter is above 0.5. The parameter is relatively stable across areas and seasons. The standard deviation of the observational shocks is on average 0.271 with standard deviation of 0.119. Because the mean log sales share difference from the outside option is -5.892, the size of the shocks is small. It indicates that learning the demand parameters defined as above will enable workers to have a precise prediction of sales. The coefficients on the temperature are all positive for chilled products and negative for hot products at the mean and at the first and third quartiles. The value is relatively stable across product categories, and areas and seasons. Table A7 displays the summary statistics of 10 products that are at 10 to 100 percentiles in the belief mean. The

values imply that there is a large variation across products and across areas and seasons within a product. This shows the importance of efficient experimentation to learn the local demand and the efficient implementation of the optimal product assortments in this retail business.

4 Algorithm for Experimentation

In this section, I propose a tractable algorithm to dynamically optimize the product assortment decisions. Then, I predict the impact of the algorithm on the basis of the estimated demand model by simulations. I show that the algorithm outperforms workers' decisions at least in the simulation.

4.1 *Ex Ante* Optimal Product Assortment Decisions

Assume the demand function is specified as in Section 3. Let $\mathbf{d}_{ikt}$ be the vector of product assortments at vending machine k in area-season i in week t . Then, the revenue of product j at the vending machine under the product assortments is:

$$r_{ijkt} = p_j \frac{\exp(v_{ijt}/\gamma_i) d_{ijkt}}{\sum_{l \in J_g} \exp(v_{ilt}/\gamma_i) d_{ilkt}} \frac{\left(\sum_{l \in J_g} \exp(v_{ilt}/\gamma_i) d_{ilkt} \right)^{\gamma_i}}{1 + \sum_{m \in \mathcal{G}} \left(\sum_{l \in J_m} \exp(v_{ilt}/\gamma_i) d_{ilkt} \right)^{\gamma_i}} \exp(\sigma_i \epsilon_{ijkt}) \quad (1)$$

$$\equiv r_{ijt}(\mathbf{d}_{ikt}, \boldsymbol{\xi}_i, \boldsymbol{\beta}_i, \gamma_i) \exp(\sigma_i \epsilon_{ijkt}),$$

where the ijt index of the mean revenue function r represents the dependency on the price and temperature in the area-season during the week. The total revenue from an area-season in a week is:

$$r_{it} = \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{J}} r_{ijt}(\mathbf{d}_{ikt}, \boldsymbol{\xi}_i, \boldsymbol{\beta}_i, \gamma_i) \exp(\sigma_i \epsilon_{ijkt}).$$

I design an agent that sequentially decides product assortments $\mathbf{d}_{it}$ while learning the demand from the past sales data $\mathbf{r}_{i,1:t-1} = (r_{i1}, \dots, r_{i,t-1})'$. Let $F_{i,t+1}$ be the belief of the agent over the local demand parameters at the end of period t , and let $B(F_{i,t+1}; F_{it}, \mathbf{d}_{it})$ be the distribution of the end-of-period belief $F_{i,t+1}$ conditional on the beginning-of-period belief F_{it} and the assortment decision in period t under the assumption that the agent is Bayesian, which I explicitly characterize below. The randomness of the belief process comes from the realization of idiosyncratic sales shocks ϵ_{it} .

Let $\mathbf{d}_{it}(F_{it})$ be a policy of the agent that maps from the belief at the beginning of the period to the product assortment for period for $t = 1, \dots, T$. Under this policy, the expected sales from the area is:

$$\begin{aligned}
& \mathbb{E} \left\{ \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{J}} r_{it}[\mathbf{d}_{ikt}(F_{it}), \boldsymbol{\xi}_i, \boldsymbol{\beta}_i, \gamma_i] \exp(\sigma_i \epsilon_{ijk}) \middle| F_{i1} \right\} \\
&= \exp\left(\frac{\sigma_i^2}{2}\right) \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{J}} \int \cdots \int r_{it}[\mathbf{d}_{ikt}(F_{it}), \boldsymbol{\xi}_i, \boldsymbol{\beta}_i, \gamma_i] dF_{it}(\boldsymbol{\xi}_i, \boldsymbol{\beta}_i, \gamma_i) \prod_{l=1}^{t-1} dB[F_{i,l+1}; F_{il}, \mathbf{d}_{il}(F_{il})] \\
&\propto \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{J}} \int \cdots \int r_{it}[\mathbf{d}_{ikt}(F_{it}), \boldsymbol{\xi}_i, \boldsymbol{\beta}_i, \gamma_i] dF_{it}(\boldsymbol{\xi}_i, \boldsymbol{\beta}_i, \gamma_i) \prod_{l=1}^{t-1} dB[F_{i,l+1}; F_{il}, \mathbf{d}_{il}(F_{il})]
\end{aligned} \tag{2}$$

The goal of the agent is to find a policy that maximizes the expected sales subject to institutional restrictions on feasible product assortments. In particular, product assortment decisions should satisfy two institutional restrictions. First, the total number of slots assigned to each product group at a vending machine is fixed at c_{igkt} . Second, a product should be filled if commanded by the centralized product assortment decision policy. Let $\underline{d}_{ijkt} = 1$ if product j is commanded at vending machine k in week t and 0 otherwise.

This is a standard Markov decision process if I regard the belief F_{it} to be the state variable. Therefore, the optimal decision rule can be determined by solving the Bellman equation:

$$V_{it}(F_{it}) = \max_{\mathbf{d}_{it}} \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{J}} \int r_{ijt}[\mathbf{d}_{ikt}(F_{it}), \boldsymbol{\xi}_i, \boldsymbol{\beta}_i, \gamma_i] dF_{it}(\boldsymbol{\xi}_i, \boldsymbol{\beta}_i, \gamma_i) + \int V_{i,t+1}(F_{i,t+1}) dB(F_{i,t+1}; F_{it}, \mathbf{d}_{it}).$$

subject to:

$$\begin{aligned}
& \sum_{j \in \mathcal{J}_g} d_{ijkt}(F_{it}) = c_{igkt}, k \in \mathcal{K}, g \in \mathcal{G}, \\
& d_{ijkt} \geq \underline{d}_{ijkt}(F_{it}), j \in \mathcal{J}, k \in \mathcal{K}.
\end{aligned}$$

for $t \in \mathcal{T}$ with a terminal condition $V_{|T|} = 0$.

I call the solution to this problem the *ex ante* optimal product assortment decisions, given the information available up to the timing of the decision. The key trade-off in this problem is between exploitation and exploration. Although the agent wants to choose assortments that are expected to yield immediate high reward, it also wants to choose assortments that help reduce uncertainty so that it can make better decisions in the future.

4.2 Decision with TS

It is not straightforward to solve this problem because of the non-regularity of the setting. First, the dimensionality of state space is extremely large because there are hundreds of products and beliefs are defined over those products. This makes it impossible to solve the problem using backward induction with brute force. Second, Gittin’s rule Gittins and Jones (1974) is not optimal because this is a combinatorial multi-armed bandit problem. Third, there are multiple vending machines in an area, and there is information spillover across vending machines as product assortment is determined for each. Caro and Gallien (2007) and Rusmevichientong et al. (2010) studied the dynamic product assortment problem and proposed heuristic policies. However, their policies are not directly applicable because they did not consider a situation where there are multiple sites to determine the product assortment. Additionally, there was information spillover. Caro and Gallien (2007) policy exploits a specific form of demand that has no substitution across products. Thus, is not empirically appropriate.

Therefore, I borrow the concept of TS, which has recently attracted considerable attention in the field of the multi-armed bandit problem. The idea is assuming a simple prior distribution on the parameters of the reward distribution of every alternative. Then, at any time step, we play an alternative according to its posterior probability of being the best. Several studies, including Scott (2010), Chapelle and Li (2011), May et al. (2012) have demonstrated the empirical efficacy of TS using simulations. Agrawal and Goyal (2012) showed that the TS achieves logarithmic expected regret for the stochastic multi-armed bandit problem. After the experiment was conducted, there was a development in the literature. The algorithm I develop below can be considered as a special case of the TS algorithm for generalized linear context bandit problems in Abeille and Lazaric (2017). According to their result, the algorithm below has a regret bound of rate $O(\sqrt{T})$.

TS works as follows. i) At the beginning of period t , for each vending machine k , the agent draws a sample from his belief F_{it} , and lets it be $\xi_{ik}^*, \beta_{ik}^*, \gamma_{ik}^*$. ii) Then, at vending machine k , the agent chooses a product assortment that solves:

$$\max_{\mathbf{d}_{ikt}} \sum_{j \in \mathcal{J}} r_{ijt}(\mathbf{d}_{ikt}, \xi_{ik}^*, \beta_{ik}^*, \gamma_{ik}^*),$$

from a feasible set of assortment. This is the optimal product assortment when the true demand parameter is $\xi_{ik}^*, \beta_{ik}^*, \gamma_{ik}^*$, for $k \in \mathcal{K}$. iii) The agent updates its belief on the basis of the sales r_{ijkt} , for $j \in \mathcal{J}$ and $k \in \mathcal{K}$. iv) It continues this process until the end of the season. Thus, the agent has product assortment plans that are optimal at every state of the world,

and it mixes it with a belief about the state of the world. There is growing literature about algorithms that solve this kind of combinatorial optimization problem when the demand parameters are known. With the nested multinomial logit choice assumption above, I can utilize the algorithm of Gallego and Topaloglu (2014) to solve a static product assortment problem with a group-by-group slot restriction.

A product assortment is likely to be chosen under the TS algorithm if it has high expected immediate rewards or high uncertainty about the immediate reward. The first factor corresponds to exploitation motivation and the second factor corresponds to exploration motivation. As the number of trials increases, the uncertainty decreases, and exploitation motivation dominates. Thus, it resolves the exploitation-exploration trade-off in an intuitive manner. One drawback of this algorithm is that it is not a forward looking decision rule. When an increased exploration value is expected, forward looking decision maker increases the degree of exploration, whereas the decision makers following the TS algorithm do not.

At the end, the choice probability at period t at vending machine k given belief F_{it} is:

$$p_{it}(\mathbf{d}_{ikt}|F_{it}) = \int 1\{\mathbf{d}_{ikt} = \operatorname{argmax}_{\mathbf{d}} \sum_{j \in \mathcal{J}} r_{ijt}(\mathbf{d}, \boldsymbol{\xi}_{ik}^*, \boldsymbol{\beta}_{ik}^*, \gamma_{ik}^*)\} dF_{it}(\boldsymbol{\xi}_{ik}^*, \boldsymbol{\beta}_{ik}^*, \gamma_{ik}^*),$$

which completes the algorithm's description. Given the choice probability $p_{it}(\mathbf{d}_{ikt}|F_{it})$, I can sample decisions of the agent, and given the belief-updating law $B(F_{i,t+1}; F_{it}, \mathbf{d}_{it})$, I can sample the belief during the next period. These are *TS* product assortment decisions.

In the next section, I conduct a field experiment introducing TS product assortment decisions to the company. To do so, I put additional practical restrictions on the model.

First, during a season, only the product-specific components of the demand function, $\boldsymbol{\xi}_i$, are learned online during the season. The group-specific components $\boldsymbol{\beta}_i$, γ_i , and σ_i are not learned. Instead, I use the average of the estimates across seasons in the local area. This provides numerical stability for the algorithm. Second, I allow the initial belief of the algorithm, F_{i1} , and the sales shock ϵ_{ijkt} , be normal distributions. Then, the beliefs, F_{it} , are all normal distributions. Third, the initial belief for $\boldsymbol{\xi}_i$ for the products that have been available since the last season is set at the *ex post* belief for them in the last season of the local area. As for the initial belief for $\boldsymbol{\xi}_i$ for new products, the belief mean is set to the mean of the belief mean of the product group, and the belief standard deviation is set to the standard deviation of the belief means of the product group. I change this part in one of the treatment groups during the field experiment.

Under this assumption, the measurement model for a product, where $d_{ijkt} = 1$, has a

linear-normal form in the unknown parameters as:

$$y_{ijkt} \equiv \ln(s_{ijkt}/s_{i0kt}) - \beta_{it} \times Temperature_{it} - (1 - \gamma_{ig})s_{ij|gkt} = \xi_{ij} + \sigma_i \epsilon_{ijkt}.$$

Therefore, the belief updating function, $B(F_{i,t+1}; F_{it}, \mathbf{d}_{it})$, for ξ_i is derived by applying a Kalman filter.

4.3 Comparing with Actual Product Assortments

I call the optimal product assortments when the demand function is evaluated at *ex post* belief of the demand parameters the *ex post* optimal product assortments. Given product assortment decisions d_{ijkt} and *ex post* optimal product assortment decisions $d_{ijkt}^{Optimal}$, the proportion of delegated slots that agree on the assortments is defined by:

$$\rho_{it} = \frac{\sum_{j \in \mathcal{J}, k \in \mathcal{K}} (1 - d_{ijkt}) d_{ijkt}^{Optimal} d_{ijkt}}{\sum_{j \in \mathcal{J}, k \in \mathcal{K}} (1 - d_{ijkt}) d_{ijkt}^{Optimal}}.$$

This is called the *success rate* of the product assortment decisions in the week. I compute this distance for every week across areas and seasons for actual and TS product assortment decisions and denote them by ρ_{it}^{Actual} and ρ_{it}^{TS} . Then, I estimate the mean distances $\mathbb{E}\{\rho_{it}^{Actual}\}$ and $\mathbb{E}\{\rho_{it}^{TS}\}$, and their 99% confidence intervals.

Table A8 and Figure A2 show the mean success rates of TSs and actual product assortment decisions and their 99% confidence intervals. Mean success rates are always higher under TS product assortments than under actual. During the first week, the mean success rates are 0.367 and 0.29 under TS and actual product assortments, and in the last week, they are 0.307 and 0.225. To formally test the null hypothesis, in which the success rates of actual product assortments are higher than those of *ex post* optimal product assortments in each week, I estimate the following linear model: For $\tau = Actual, TS$,

$$\rho_{it}^{\tau} = \sum_{l \in \mathcal{T}} \{\rho_{0l} + \rho_{1l} 1\{\tau = Actual\}\} 1\{l = t\} + \epsilon_{it}^{\tau},$$

with and without the inequality restrictions $H_0: \rho_{1l} \geq 0$ for $l \in \mathcal{T}$, where ϵ_{it}^{τ} is assumed to be an i.i.d. mean zero random variable. I estimate the model with and without the inequality restrictions, compute the test statistics, and estimate the critical value using a bootstrap of resampling the whole data with replacements for 1,000 times. The resulting test statistic is 20181.48, and the p-value is 0.00. Hence, the null hypothesis that the success rates of TS product assortments is lower is rejected.

I also compare the regret under TS to actual product assortments. In the current context, the absolute regret of a product assortment policy in a period is defined as the expected revenue under the *ex post* optimal product assortments minus the expected revenue under the product assortment policy during the period. Table 4 and Figure 5 show the mean and 99% confidence intervals of the regret under TS and actual product assortments over time relative to the expected revenue under the *ex post* optimal product assortments. The regret is always lower under TS than under actual product assortments. To formally test the null hypothesis that the regret of TS product assortments is higher than those of actual product assortments in each week, I conduct the same test used for the success rates. The resulting test statistic is 24302.38, and the p-value is 0.00. Hence, the null hypothesis is rejected. During the first week, the regrets are 0.114 and 0.155 under TS and actual product assortments, and 0.028 and 0.069 during the last week. Thus, TS product assortment decisions are expected to yield higher revenue than the actual product assortments under the assumption that the demand function is correctly specified.

5 Field Experiment

5.1 Experiment Design

The unit of assignment is an area, and the sample cover the areas in Tokyo. The experiment is conducted from September 2016 to February 2017. I randomly divide the areas into three groups: control, treatment I, and treatment II. In the control group, I do not intervene at all, letting the workers decide product assortments as usual. I first stratify the areas per the responsible subcontractors and then assign the treatment status within each stratum. In treatment I, I provide the instructions generated from the TS algorithm that starts from the prior belief, which *does not* incorporate the workers' relative sales prediction for the new products. This is described in Section 5.2. In treatment II, however, I provide instructions generated from the TS algorithm that starts from the prior belief, which *does* incorporate the workers' relative sales prediction for the new products. To the workers, I announce that the algorithm in treatment I does not use the workers' prediction but that the algorithm in treatment II does. The workers are also notified to which group they belong.

I then compare the correlation between the actual product assortment decisions and the TS product assortment decisions. The TS product assortment decisions with and without the workers' knowledge integration in the control group are not shown to the workers but have been computed in the background. I did this to see (i) if treatment I's decisions are more correlated than those of the control group, and (ii) if treatment II's decisions are

more correlated than treatment I's. If the correlation between the actual and TS product assortment does not differ between the control group and treatment I, it indicates that workers are not following the algorithmic advice. If the correlation is higher in treatment II than in treatment I, the integration of the workers' forecasts into the algorithm increases the conformity of workers to the algorithm.

To make the whole procedure implementable, the message generated from the algorithm is made coarse. First, the advice is only given monthly. By the third week of each month, I train the algorithm using the data available at that point and generate product assortment decisions under the 10 year average temperature at the end of the next month. I then give this information to the workers. Second, the product assortment decisions are aggregated at the area-product level and not given at the vending-machine-product level. Thus, the advice takes the form of "coca cola should reach 13% of the soda category slots in your area by the end of next month" for respective products. Because the TS algorithm is stochastic, I generated product assortments for 100 times according to the TS algorithm and took the average to obtain the target slot shares. The current slot share of each product is also provided for workers' information.

The message is summarized in a spreadsheet for each area, sent separately to the targeted areas. Figure 6 is a screenshot of the spreadsheet. I did not provide information about the other areas. However, I could not prohibit the workers from sharing the information.⁷ The slot share of each product and the sales at each area is monitored every month, and the summary of the results are sent to the workers at the same time with the algorithmic advice. These spreadsheets are distributed by the company manager.

All instructions are given through the manager of the principal company, and I do not appear to the workers. The manager promotes the algorithm in every monthly meeting with the subcontractors. But he does not force them to absolutely follow the instruction, because it is a requirement beyond the current contract. The principal company does not change the existing incentive system to encourage the subcontractors to follow the algorithm. The fee to the subcontractors is kept to be linear in the revenue. Thus, the following results should be interpreted as facts in this type of environment. How to communicate with workers to involve them and how to design incentives to promote algorithm adoption are interesting

⁷Each worker belongs to an office in which workers in the neighborhood areas regularly show up. It is possible that a worker who received the algorithmic advice shares the information with the other workers belonging to the same office. In the online appendix, I conducted the regression analyses that control for the average algorithmic advice in each office, in order to control for the potential information sharing at the office. In particular, for each worker, I calculated the average net number of targeted slots of products advised to the other workers belonging to the same office. Table A14 shows the estimation results where the average advice is controlled. This shows that the average advice is not significant and the other results remain the same.

issues but left for future research.

Because it is expected to be hard for workers to fully understand the algorithm, I briefly explain in a document that I use the sales data of each area since 2013 until just before the experiment, to estimate the popularity of each product, and that I update the estimate using the latest sales data each month, considering the sensitivity of the sales to temperatures and the composition of the products in the vending machines. Then, I explain that the algorithm finds optimal product assortment decisions on the basis of these estimates so that it balances exploration and exploitation.

5.2 Workers' Belief for New Products

In August, the principal company assembles the representatives of the subcontractors and provide them with detailed information about new products and new policies in the Fall/Winter season. The representatives of the subcontractors then bring back the information to their offices and share with their workers. In this way, the subcontractors and the workers can know all the coming products in advance. I asked the principal company to conduct a survey asking workers their forecast for the products planned to be introduced during the coming season, Fall/Winter 2016, as one of the items on the agenda of this meeting. During this season, 24 distinct products were introduced. Some were introduced in early fall and some in late fall.

In the forecasting survey, I showed the workers a list of new products that displayed the name and product characteristics, including an image of the package. Then, I asked the workers to report the predicted sales volume relative to the categorical average in their area during a given week under a hypothetical temperature. For example, suppose that a worker predicts the sales of a new product called *Green* in his area in the week starting October 17. Suppose that *Green* belongs to the green tea category. It is explained that the average temperature in the week starting October 17 is 16°C. Suppose that the worker believes that products in the green tea category will sell on average 100 units whereas *Green* will sell 120 units in the week. Then, he reports $1.2 = 120/100$ as a score of *Green*. The survey format copies the format the principal company and subcontractors regularly use to review the sales of a product. They regularly monitor the relative sales of a product to the category average and hence have some sense of this metric. For the products that are released by the week starting October 17, I asked the workers to predict the sales during the week when the temperature is 16°C. For the products released by the week starting November 21, the workers were advised to predict the sales during the week when the temperature is 8°C. For the other products that are all released by the week starting December 19, I asked the

workers to assume that the temperature was 5°C. These temperatures were based on the 10 year average in the Tokyo metropolitan area.

These tasks are demanded by the company as a part of their job. There is no monetary reward specific to the tasks. However, I announced them so that the results would be publicly shared among workers, and the top 3 workers would be identified at the end of the season. I did not employ a sophisticated method to elicit personal beliefs because of managerial resource constraints.

The summary of the prior belief for new products by workers is reported in Table 5, according to the release periods and the relevant forecasting setting. On average, predicted sales was less than 1.0, implying that workers tend to undervalue new products relative to existing products. Some workers predicted extreme sales, such as 0.1 and 5.8.

To determine whether their estimate was informative, I computed *implied* ξ_{ij} , denoted as $\xi_{ij}^{implied}$, given the relative sales prediction from the workers. I used the *ex post* belief of the demand parameters as from the end of Spring/Summer season 2016, to compute the predicted sales of all available products, except for the new products under each setting. Second, I multiplied the workers' relative sales prediction to the category average sales to obtain the predicted sales of the new products under each setting given the prior belief of the workers. Finally, I computed $\xi_{ij}^{implied}$ by inverting the within-category choice probabilities.

On the other hand, I took within-category average of ξ_{ij} s in the *ex post* belief from the end of the Spring/Summer season 2016, denoted by $\bar{\xi}_{ig}$. This is a least possible guess for the demand parameters of new products, using only sales and product characteristics information up to the beginning of the Fall/Winter season 2016. One can add other product characteristics such as brands and beverage companies. However, in a season, there are at most around 10 products in a category. Because the sample size is small, adding characteristics to estimate and predict the demand for new products can make the prediction inaccurate. Moreover, it becomes harder to explain how the algorithm works to the workers. For these reasons, I simply took the latest within-category average to form the initial belief for the new products.

To see how these values, $\xi_{ij}^{implied}$ and $\bar{\xi}_{ig}$, are informative, I regressed the *ex post* belief of ξ_{ij} from the end of the Fall/Winter season on these values. The *ex post* estimate of ξ_{ij} , denoted by $\hat{\xi}_{ij}^{expost}$, is obtained in the same way as that the Section 3.1. Table 6 shows the regression results, when only $\bar{\xi}_{ig}$ is included, only $\xi_{ij}^{implied}$ is included, and when both $\xi_{ij}^{implied}$ and $\bar{\xi}_{ig}$ are included. All coefficients except for the constant, are positive and statistically significant. However, the size of the coefficient is larger for $\bar{\xi}_{ig}$ than for $\xi_{ij}^{implied}$, and the model fit in terms of R^2 is also higher for model (1) than for model (2). In model (3), where $\xi_{ij}^{implied}$ and $\bar{\xi}_{ig}$ are included, the size of the coefficient for $\xi_{ij}^{implied}$ is almost zero, and the model

fit does not improve compared with model (1). This result indicates that the workers may have some prior knowledge about the new products but they are not particularly informative once the company exploits the sales and product characteristics information available up to decision time.

5.3 Average Response

To study the effects of providing the algorithmic advice, I estimate the following model:

$$\begin{aligned}
 ActualNumSlot_{ijt} = & \alpha_1 NumNetTargetedSlot_{ijt} \\
 & + \alpha_2 NumNetTargetedSlot_{ijt} \times Advice_i \\
 & + \alpha_3 NumNetTargetedSlot_{ijt} \times Integration_i \\
 & + \alpha_4 NumCentralizedSlot_{ijt} \\
 & + Controls + \epsilon_{ijt},
 \end{aligned} \tag{3}$$

where i is an area, j is a product, t is a target week, $ActualNumSlot_{ijt}$ is the number of actual slots devoted to the product in the area during the target week, $NumCentralizedSlot_{ijt}$ is the number of slots required by the centralized product assortment policy, and $NumNetTargetedSlot_{ijt}$ is the number of additional slots suggested to be filled by the algorithms. $Advice_i$ is the dummy of both treatment groups, for which workers were informed of the output of the TS algorithm. $Integration_i$ is the dummy of treatment II group. $NumNetTargetedSlot_{ijt}$ for the control group is based on the algorithm that does not utilize the workers' prior beliefs. The main parameters of interest are α_2 and α_3 , which capture the difference in correlation with the algorithmic advice between treatment statuses. I estimated model controlling for product-specific and year-week-specific fixed effects for a robustness check.

Table 7 shows the estimation results. (i) $Advice_i$ did not significantly affect the workers' product assortment decision, and (ii) the integration of workers' belief statistically significantly increased the conformity of workers with the algorithmic advice. This implies that (i) there was a barrier to adopting the algorithm and that (ii) integrating the workers into the algorithm design process reduced the barrier to some degree. The partial correlation between the actual product assortment decisions and the TS product assortment decisions, in terms of the product-level number of slots, was 0.015 higher in treatment II compared to the control group. The correlation, however, was 0.118 in the control group. The change in the extent of correlation can be regarded as large enough relative to the baseline correlation (13.011%) and as the effect of just integrating the workers' prior belief into the algorithm.

The difference in the correlation between treatment I and II may have been caused by two

channels. First, the pure psychological effects of integrating workers into the algorithm design process may have caused the difference. Second, the instructions provided in treatment I were more similar to the workers’ own idea, perhaps influencing the effects. The design of the field experiment does not fully distinguish between these two scenarios. However, there was little difference in the product assortment decisions between algorithms with and without the workers’ prior belief. The correlation between these two product assortment decisions, in terms of the number of targeted slots is 0.979. This is because the effects of the prior belief quickly diminish once actual sales are observed. Therefore, the first channel will likely dominate.

The next question is whether the changes in the product assortment decisions were large enough to cause changes in the revenue. To study this issue, I estimate the following model:

$$Outcome_{it} = \alpha_0 + \alpha_1 Advice_i + \alpha_2 Integration_i + Controls + \epsilon_{it}, \quad (4)$$

where $Outcome_{it}$ is either sales volume or value of month t in area i where the mean sales volume or value of the area in the last year in the same season is subtracted and divided by the standard deviation of the sales volume or value. Table 8 summarizes the results. The dependent variable for model (1) and (2) is sales volume normalized by the last year’s average sales volume in the area. The dependent variable for model (3) and (4) is sales value normalized by the last year’s average sales value in the area. Model (2) and (4) control for month-specific fixed effects. The results show that the treatments do not have statistically significant effect on the sales. Even though the coefficients of $Integration_i$ are positive, because the coefficients on $Advice_i$ are negative and larger, the sales are overall slightly lower among the treated areas. These results are similar to the other outcome metrics, such as sales level, total sales during the experiment period, and year-on-year monthly sales growth. In the next section, I only show the results of the normalized sales volume, because the results with other outcome variables are similar.

The results of the average responses so far are not rosy. They show that workers are on average reluctant to follow the algorithm. Even though integrating workers into the design process can mitigate this problem, the impact is still not large enough to cause significant changes in the revenue. However, different types of workers may respond differently to the introduction of an algorithm. If this is true, we may be still able to improve performance by carefully choosing subsets of workers who are more favorable to using an algorithm or by changing the communication with workers. To explore this possibility and to better understand when and why workers follow an algorithm, I next study the heterogeneous responses of the workers.

5.4 Heterogeneous Responses

The workers' reactions are expected to be heterogeneous, depending on the type of the worker, such as experience, confidence with one's product assortment decisions, and trust in the algorithm. Moreover, workers may selectively change their behaviors across vending machines and products. For example, a worker may follow an algorithm only on low-traffic vending machines, fearing the risk of changing behaviors on high-traffic vending machines. For the same reason, they may avoid expensive products when they experiment with the algorithmic advice. To facilitate the interpretation, in the following analyses, all of the continuous variables except for the number of slots are normalized by subtracting the mean and dividing by the standard deviation. Because the following analyses extensively use non-experimental variations, the results should be only considered as suggestive evidence and should be interpreted with reservation. Because the results are almost the same across the way of controlling for the unobserved fixed effects, I only report the last results and report others in the online appendix.

Area/worker-level heterogeneity Firstly, I study how the performance of a worker is related to the response. One of the performance metrics that is often used in the context of a bandit problem is a regret. I compute the difference in the weekly regret of the actual assortment decisions and the TS product assortment decisions of each area during the Spring/Summer 2016. I took the average of the weekly regret for each area and used this measure, $Regret_i$, as the worker's performance in an area compared to the TS algorithm in the last season. Although the identity of workers in an area can change across seasons, the change is not substantial within the same year, according to the company's managers. These variables are interacted with $NumNetTargetedSlot_{ijt}$, $NumNetTargetedSlot_{ijt} \times Advice_i$, and $NumNetTargetedSlot_{ijt} \times Integration_i$. Column (1) of Table 9 shows the estimation results when interacting $Regret_i$ with the variables in the previous model. The result shows that if the regret is higher i) workers follow the algorithmic advice more, and ii) they become reluctant when their belief is integrated into the algorithm. This result is intuitive: a low performer prefers a plain algorithmic advice and a high performer prefers to influence the algorithm.

The estimates show that correlation between algorithmic and actual assortments are zero for workers with an average level of regret, both with and without integration. When the regret of a worker is at the level of one-standard-deviation, the correlation increases by 0.046 without integration, whereas with integration it additionally decreases by 0.064. This is one of the reasons for the mixed results on the average treatment effects. A low performer conforms to the algorithmic advice if his opinions are integrated, but a high performer resists

it. On the other hand, a lower performer resists the algorithmic advice if their opinions are integrated, but a high performer conforms to it. Thus, the average impacts can be either positive or negative depending on the distribution of low and high performers.

A worker can have a higher performance either because the worker is able or the area is easy to operate. For example, if the sales volatility is high, then it would be harder to learn an area’s demand. To distinguish these channels, secondly, I study the heterogeneity due to sales volatility across areas. I compute the standard deviation of weekly sales volume of a product in an area in the 2016 Spring/Summer season. This variable is also interacted with $NumNetTargetedSlot_{ijt}$, $NumNetTargetedSlot_{ijt} \times Advice_i$, and $NumNetTargetedSlot_{ijt} \times Integration_i$. Column (2) of Table 9 summarizes the estimation results. The results show that even after controlling for the sales volatility in an area, the degree of regret or performance affects the effectiveness of the algorithmic advice in the same direction. The result shows that a worker in an area with higher sales volatility is less likely to follow the algorithm without integration. However, with integration, a worker in an area with higher sales volatility is more likely to follow the algorithm. Note that the coefficients of the interaction between the area sales volatility and the net number of targeted slots in the control group is positive and significant after controlling for product- and year-week-specific unobserved fixed effects. This means that in such an area, the worker’s product assortment decisions are already closer to the TS algorithm. One possible interpretation is that workers in a volatile area have already made an effort in learning the local demand and are confident in their decisions. Thus, they are less likely to follow the algorithmic advice that does not integrate their opinions and are more likely to follow one that does integrate their opinions.

Third, I study how work experience influences the worker’s response. To see this, I construct $Experience_i$ dummy; it takes a value of 1 if the work experience is more than 5 years and 0 otherwise. Column (3) of Table 9 shows that experienced workers more strongly prefer an algorithm in which their opinions are integrated by 0.032 points. However, interestingly, experience is irrelevant in the correlation between the actual and the net number of targeted slots in the control group: the coefficients are small and not statistically significant. Thus, work experience only changes the workers’ attitudes toward the algorithmic advice and whether their opinions to be integrated, but does not seem to change the quality of their product assortment decisions.

Fourth, I investigate the possibility that the degree of delegation to a worker may affect a worker’s compliance to two types of algorithmic advice. As I described in Section 2, workers are responsible for about 30% of the slots in the vending machine. However, the actual share of delegated slots is slightly different across vending machine types and the aggregate share of delegated slots in an area also varies across areas because different areas

have different combinations of vending machines. If a worker is given more delegated slots, then that worker is able to more flexibly react to the algorithmic advice. To examine this, I construct $Delegation_{it}$, the share of delegated slots in an area during the target month. This is based a nominal number of slots to be delegated in the season and not the actual number of slots filled with products not required by the principal company. Column (4) of Table 9 summarizes the estimation results. The results show that the degree of delegation affects the worker’s response to the algorithmic advice in a similar manner as sales volatility. A worker in an area with more delegated slots is less likely to follow the algorithm without integration. However, integration makes a worker in an area with more delegated slots more likely to follow the algorithm. However, there is a difference in the effect on the control group. In the control group, the coefficients on the interaction between the degree of delegation and the net number of targeted slots is negative and significant after controlling for product- and year-week-specific unobserved fixed effects. Two interpretations seem to be available. First, similar to the case with sales volatility, a worker in more delegated area might have already put some effort in learning the demand. Second, higher degree of delegation may imply higher operational cost to follow the algorithm.

In sum, a worker is more likely to follow the plain algorithmic advice if he has higher regret compared to the TS algorithm in the last season, faces lower sales volatility, or has a lower degree of delegation. Work experience is not significantly correlated with the compliance with the plain algorithmic advice. On the other hand, integrating a worker’s opinion makes the worker more likely to follow the algorithmic advice if the worker has lower regret compared to the TS algorithm in the last season, faces higher sales volatility, had more work experience, or has higher degree of delegation.

The heterogeneity in worker response leads to the heterogeneity in the effects on the sales. To see this, I interact the previous area/worker-level variables with $Advice_i$ and $Integration_i$ in the sales volume regression. Table 10 summarizes the results for the models controlling for month-specific unobserved fixed effects. The signs of the coefficients are largely in line with the signs of the corresponding coefficients in the product assortment regressions, although statistical significance is lost in some cases.

If a worker has higher regret, the worker is more compliant to the algorithmic advice without integration. The result shows that the normalized sales volume is also higher, although it is only statistically significant before the degree of delegation is controlled. On the other hand, if worker regret is high, the worker is less compliant to the algorithmic advice with integration. This result shows that the sales are also lower, although they are not statistically significant. If the sales are volatile, the worker is less compliant to the algorithm without integration and more to the one with integration. The signs in column

(4) have the same patterns, but the effects are weak. The effects, however, are evident for work experience and degree of delegation. If a worker was more experienced, the worker was more compliant to the algorithm with integration, whereas not more or less compliant to the algorithm without integration. The normalized sales is also statistically significantly higher when the worker is experienced and revealed to the algorithm with integration, whereas there is no such an effect with the algorithm without integration. If a worker was given more delegated slots, the worker is less compliant with an algorithm without integration and more compliant to an algorithm with integration. The signs for the normalized sales regression have the same signs and are statistically significant. Thus, the algorithmic advice affects worker behavior, and by doing so, improves the performance. However, the effects are subtle and heterogeneous, which resulted in mixed average impacts.

Machine-level heterogeneity So far this study has regarded the average response of a worker in his area. The next step is to study heterogeneity across vending machines in an area. First, I classify vending machines into high and low sales volume groups and then into high and low sales volatility groups based on the sales volume in 2016 Spring/Summer season. Next, I compute the actual number of slots, net number of targeted slots, and number of centralized slots respectively for vending machines belonging to each group for each area. Now script i refers to the group that pairs the area and the vending machine group. The variable $HighSalesVolumeMachine_i$ takes a value of 1 if the average sales volume of the machine is in the upper half and 0 otherwise. The variable $HighSalesVolatilityMachine_i$ takes a value of 1 if the sales standard deviation of the machine is in the upper half and 0 otherwise.

Column (5) of Table 9 displays the results. First, the signs and statistical significance of the existing variables are unchanged. Second, the results show that a worker is more reluctant to follow the algorithm when stocking machines with high sales volume. A worker may not want to risk experimenting with a new assortment policy in high traffic machines. Third, the results show that a worker is more willing to follow the algorithm when stocking machines with high sales volatility. A worker may find it difficult to adapt to the demand of such a machine and would rather rely on the algorithm. Fourth, integration is more effective in machines with high sales volume and low sales volatility. Workers have more confidence in their decisions when stocking these types of machines: therefore, they need more justification to convince them to follow the algorithm. Finally, on machines with high sales volatility a worker's product assortment decisions are less correlated with the TS algorithm, even though they are closer in an area with high sales volatility. One interpretation is that a worker puts more effort into learning the local demand in a volatile area but he still fails to find an

assortment that is similar to the TS algorithm in a volatile machine. Another interpretation is that there is a machine-level heterogeneity in the demand within an area and the worker can successfully adapt to the local demand, whereas the TS algorithm in this study, which only allows for area-level heterogeneity, fails to adapt to it. If the latter case is true, then workers should be less willing to follow the algorithm in machines with high volatility. However, as shown above, workers are more willing to follow the algorithm in machines with high volatility. Therefore, the former case seems to be true in the current context.

The column (5) of Table 10 shows the estimation results for effects on normalized sales with machine-level heterogeneity. As I include the vending machine group variables, some of the statistical significance of the existing variables' coefficients disappear. The signs of the effects of vending machine groups on sales are the same as the signs of the effects on assortment decisions. The results show that the effects of algorithmic advice on sales is lower in machines with high sales volume, because the effects on product assortment decisions was weaker. Integrating workers' opinions has a higher impacts on sales for machines with high sales volume because it has a higher impacts on product assortment decisions as well.

5.5 Difference in Actual and TS Product Assortments

The analysis so far revealed the conditions under which workers are willing and unwilling to follow an algorithm. The next sections explores how the actual product assortments and TS product assortments are different. To see this, I estimate the following model:

$$\begin{aligned}
ShareDifference_{ijt} = & New_j + PrivateBrand_j + Hot_j \\
& + AluminumBottle_j + AluminumCan_j \\
& + Volume_j + Price_j + CocaByCoca_{ij} + Coca_j \\
& + (New_j + PrivateBrand_j + Hot_j \\
& + AluminumBottle_j + AluminumCan_j \\
& + Volume_j + Price_j + CocaByCoca_{ij} + Coca_j) \times Advice_i \\
& + (New_j + PrivateBrand_j + Hot_j \\
& + AluminumBottle_j + AluminumCan_j \\
& + Volume_j + Price_j + CocaByCoca_{ij} + Coca_j) \times Integration_i \\
& + Controls + \epsilon_{ijt},
\end{aligned} \tag{5}$$

where $ShareDifference_{ijt}$ is the difference in the percentage shares of products in the actual and TS product assortments, New_j is a dummy for new products, $PrivateBrand_j$ is a dummy for the principal company's private brand, Hot_j is a dummy for hot beverages,

and *AluminumBottle_j* and *AluminumCan_j* are dummies for aluminum bottles and cans, respectively: if both are 0 it represents a plastic bottle. *Volume_j* is the standardized volume of the package, and *Price_j* is the standardized price. *Coca_j* is a dummy for products produced by the Coca Cola Company. Because there are subcontractors who are affiliated with the Coca Cola Company, I include a dummy for coca cola products served by subcontractors affiliated with the Coca Cola company. I control for year-week-specific unobserved fixed effects. This regression reveals the difference in the product assortment shares in the control group and how the provision of algorithmic advice and integrating opinions correct or possibly enlarge the difference.

Table 11 summarizes the estimation result. The result shows that, on average, the actual product share is 0.23% points lower than the TS algorithm share for a default product. Workers tend to assign 0.08% points more share to new products, 0.27% points more to the private brand, 0.29% points more to hot beverages, 0.48% more to aluminum bottles and 0.18% points less to aluminium cans. Workers assign 0.032% points more shares to a one-standard-deviation large package and 0.049% points less shares to a one-standard-deviation expensive beverage. Interestingly, workers tend to assign too many shares by 0.22% points to the Coca Cola products, whereas subcontractors who are actually affiliated with the Coca Cola Company rather tend to suppress the share by 0.1% points. It will be worth discussing other patterns. Workers' overstocking new products may be because they have higher uncertainty about new products' demand than the algorithm. Workers may expand shelves for hot beverages too early because they do not want to procrastinate this tedious task. Overstocking bigger beverages and aluminium bottles and understocking aluminium can may be because workers tend to follow their rule-of-thumb regarding the composition of the package types. Understocking pricey beverage is also understandable. Selling pricey but high value-added beverages is relatively a new practice in the beverage vending machine business.

The effects of the algorithmic advice on the difference in the shares are suggestive. Under the algorithmic advice, the shares of new products, products in aluminum bottles and in aluminum cans, large volume products, and expensive products, are all closer to the TS algorithm. The changes are statistically significant. The share of Coca Cola products and those products served by Coca Cola affiliated subcontractors are also statistically significantly getting closer to the TS algorithm. The share of hot beverages is also closer but not statistically significant. Thus, the misallocation of slots compared to the TS assortments across product types is adjusted under algorithmic advice, even though the overall correlation was not significantly different as we saw in the previous sections. The reason is that the degree of adjustment is still too minor. Too see sizable impacts on the overall conformity

and the sales, we need much larger size of the adjustment.

Integration further adjusts the allocation of shares. The share of hot beverages is closer by 0.11% points with integration and the changes are statistically significant. The misallocation of slots to aluminum bottles is also corrected further. The difference in the impacts on sales between the algorithmic advice with and without integration seems to be driven by the allocation of slots to hot and cold beverages. This makes sense because the timing of introducing hot beverage when the season’s transition from summer to winter is one of the most important decisions that affects the sales, according to principal company’s managers. In the last several years, the company has been struggling with the optimization of this margin. However, it is not clear why integrating workers’ opinions about the demand for new products would affect the workers in this margin, whereas the effects of integration on the share difference in the new products are not statistically significant.

One exception is the shares of private brand products. The difference in the share between actual and TS assortments is wider under the algorithmic advice and even wider with integration. In this case, the workers might have adjusted the shares across product types by reducing the shares that have been assigned to the private brand products. The effects of integration on the share difference in the other margin are not statistically significant.

Finally, I computed how much workers earned if they were fully compliant to the TS product assortments. Table 12 shows the summary statistics of regret under actual and TS product assortment during the experiment period, based on the ex-post belief at the end of the experiment period. On average, the regret under the actual assortments was 12.0% relative to the ex-post optimal assortments. On the other hand, the regret was, on average, 5.3% under the TS assortment. The marginal gain of fully adopting the TS assortments was on average 6.8% of the ex post optimal revenue. The potential impact of 6.8% is reasonably large. However, according to Table 9, the conformity of workers increases at most by a few percentage points even at the 1 standard deviation away from the average in the favorable direction. Thus, to achieve this maximal gain, quite a few favorable conditions have to be satisfied at the same time. This seems to be the reason why the overall impact on the sales was limited.

5.6 Follow-up Survey

The experiment revealed that there are some barriers to let workers follow an algorithm. To infer the reasons for this, I review the responses in the follow-up survey. After the experiment, I conducted a review survey at the end of March 2017. Because this survey was conducted by the company’s managers, the response rate was 100%.

The survey first asked, "Did you trust the algorithm based on the sales data *before* the experiment started?" and "Did you trust the algorithm based on the sales data after the experiment?" The answer choices were on a 4-point Likert scale: "strongly agree", "agree", "disagree" or "strongly disagree". Table A15 summarizes the answers. Most workers disagreed both before and after the experiment. However, overall, after the experiment, the proportion of workers who agreed increased, and those who disagreed decreased. Before the experiment, 36.3% disagreed and 20.1% strongly disagreed. However, after the experiment 33.0% disagreed and 13.4% strongly disagreed. Before the experiment, 32.4% agreed and 0.34% strongly agreed. However, after the experiment 40.8% agreed and 0.39% strongly agreed. This pattern is universal across the control and treatment groups, except for the proportion of "strongly agree" in treatment I.

To see if the change was statistically significant, I conducted an ordered probit, in which responses are ordered from "strongly disagree" to "strongly agree", and the regressor included an after-experiment dummy. I also included a dummy that indicated the average work experience in the area exceeded 1 year. The regression results are shown in Table A16. The first column reflects the entire sample. The after dummy is 0.42 and statistically significant. The second to fourth columns use samples from each treatment status. They show that the after dummy is statistically significant only in the control group and that the coefficients are higher for the control group. Read literally, this indicates that trust only increases in the control group where algorithmic advice is not provided. This pattern is consistent with Dietvorst et al. (2016), which found that participants, who saw the algorithm perform were less confident in it and less likely to choose it over an inferior human forecaster. They conjectured that this was because people were less tolerant of machines' errors than human errors. The fifth column uses the entire sample and includes the experience dummy. The experience dummy is negative, indicating that the experienced workers are less likely to trust the algorithm, but it is not statistically significant.

The next section of the survey was a multiple-choice question that asked the following: "What kind of improvement is needed for you to follow the algorithm?". The possible answers were the following: a reliable demand prediction, the integration of workers' predictions, the integration of managers' predictions, a detailed explanation of the algorithm, coordination with the manager's orders, a reduced burden of the work, machine-level instructions, inside-station-level instructions, station-level instructions, a mobile interface, a visual interface, coordination with the product removal schedule, and compensation for losses. The workers could select as many items as they wanted in their answers. The responses were summarized in an unweighted and a weighted manner. In the unweighted result, for each item, I counted the number of areas that included that item in the answer. In the weighted result, for each

item, if the answer in an area included the item, I divided 1 by the number of items included in the answer in the area. Then, I summed up the weighted numbers across the areas. Table 13 summarizes the result. The patterns are qualitatively the same between the weighted and unweighted summaries, and across the treatment statuses.

The most required improvement is the reduced burden of the work. Overall, the scores are 101 (unweighted) and 32.2 (weighted). This highlights the friction that may only exist in an *offline* business. In an *online* business, decisions made by an algorithm can be automatically implemented, once the code is written. However, in an *offline* business, such as a retail business with physical stores, decisions made by an algorithm have to be physically implemented. In the future, in some fields, robots may implement decisions, but today humans are needed to implement the decisions. Unless the physical implementation costs are removed, introducing algorithms into businesses cannot achieve the desired efficiency. In addition, a substantial share of workers responded that they wanted coordination in the product removal schedule. The factor is related to the workers' burden. Because beverages are seasonal products, some products are removed at a point in time. This requires workers to resolve their stock investment decisions. If the algorithm does not take this into account (the current algorithm did not), the workers' tasks can fail.

The second most required improvement is the reliability in the baseline demand prediction. The scores are 76 (unweighted) and 21.1 (weighted). The current demand model lacks several features that local workers are possibly aware of. For example, in Section 3.2, I showed that residuals are serially correlated. Moreover, there may be machine-level within-area heterogeneity in the demand parameters and the substitution across products may be more complicated than the nested-logit model.

Relatedly, a substantial proportion of workers required more detailed explanation of the algorithm. In addition to improving the algorithm, establishing better communication with workers is important. The question of how to better communicate with the workers is left for future research. Another frequently observed request is to provide instructions at the vending machine level. As explained, the optimal product assortment decisions are provided only at the area level. This was expected to *reduce* the information processing costs of the workers. However, the workers may find it easier to receive detailed instructions and then simplify the content by themselves.

The extent of the integration of workers' predictions were also substantial. The scores are 52 (unweighted) and 15.4 (weighted). Interestingly, the workers did not demand to integrate the managers' predictions. The scores are only 5 (unweighted) and 1.1 (weighted). The contrast in this case may not be between algorithm and human, but between algorithm and "me". However, in this analysis, the reliability of the managers' predictions is not controlled.

Giving workers a sense of control is one thing that can be easily improved by introducing an algorithm into business. The results from the previous section support the notion that a sense of control facilitates workers’ use of algorithm advice.

6 Conclusion

In this paper, I studied if and under what conditions workers would follow the advice of an algorithm, in the context of experimentation in product assortment decision for a beverage vending machine business. First, I developed an algorithm to automate the task of experimenting with product assortments in the current business setting. This demonstrated that the algorithm could increase revenue, at least in a simulation. The simulation exercise showed that the potential impact of automation was significant in the current business setting. In the field experiment, however, I found that the algorithm failed to increase the revenue. The reason was that the workers, on average, did not follow the algorithm. The field experiment, however, revealed that integrating workers’ forecasts into the algorithm would slightly encourage workers to follow it. Analyses that utilize non-experimental variations in the worker and the business area characteristics showed that workers’ responses to the algorithm were substantially heterogeneous. Specifically, a worker was more likely to follow the algorithm if the worker had higher regret, faced lower sales volatility, or had a lower degree of delegation. Integration was more effective if the worker on the contrary had lower regret, faced higher sales volatility, had more work experience, and had higher degree of delegation. The changes in the assortments under algorithmic advice with integration were the most evident in the relative share of slots across cold and hot beverages.

Acemoglu and Restrepo (2017) developed a framework for thinking about automation and its impact on tasks, productivity, and labor demand. According to their model, there are several countervailing channels against the displacement effect of automation. One of them is a productivity effect. In the current setting, if the company could fully automate tasks, productivity may be enhanced but further workers will be displaced. If the company compromises with partial automation, whether the productivity effect prevails depends on whether the company can find a way to resolve algorithm aversion and conflicts of interest. This paper highlighted these issues in the context of experimentation in product assortment decisions.

One limitation of this research is that it is a case study. However, this work should also offer valuable insights into the impacts of automation in the other industries. Product assortment decisions are common across consumer businesses, and the current analysis should have implications for these businesses, especially for the industries that operate in real-

world shopping sites such as supermarkets and convenience stores. This analysis did not consider joint pricing and product assortment decisions because pricing was not an issue in the current setting. It would be possible to extend this framework to study such a joint decision problem, by replacing the algorithm that obtains *ex post* optimal product assortments with one developed for joint pricing and assortment problems, such as Wang (2012).

Last, the reasons of the failure of the algorithm, and the solutions to these problems, were not extensively investigated. This is largely due to the small sample size of the current setting in terms of the unit of treatment assignment. I could only assign treatment across approximately 200 areas. This limited the number of hypotheses I could investigate in one field experiment. A greater focus on this issue could produce interesting findings that better account for the impacts of automation.

References

- Abeille, Marc and Alessandro Lazaric**, “Linear Thompson Sampling Revisited,” *Electronic Journal of Statistics*, 2017, 11, 5165–5197.
- Acemoglu, Daron and Pascual Restrepo**, “The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment,” *SSRN Electronic Journal*, 2017.
- Aghion, Philippe, Maria Paz Espinosa, and Bruno Jullien**, “Dynamic Duopoly with Learning through Market Experimentation,” *Economic Theory*, September 1993, 3 (3), 517–539.
- Agrawal, Shipra and Navin Goyal**, “Analysis of Thompson Sampling for the Multi-Armed Bandit Problem,” in “Proceedings of the Annual Conference on Learning Theory (COLT), JMLR Workshop and Conference Proceedings,” Vol. 2326 2012, pp. 1–39.
- Anand, Bharat N. and Ron Shachar**, “Advertising, the Matchmaker,” *The RAND Journal of Economics*, June 2011, 42 (2), 205–245.
- Atkin, David, Azam Chaudhry, Shamyla Chaudry, Amit K. Khandelwal, and Eric Verhoogen**, “Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan,” *Quarterly Journal of Economics*, 2017, 132 (3), 1101–1164.
- Bartling, Björn, Ernst Fehr, and Holger Herz**, “The Intrinsic Value of Decision Rights,” *Econometrica*, 2014, 82 (6), 2005–2039.
- Bergemann, Dirk and Juuso Valimäki**, “Experimentation in Markets,” *Review of Economic Studies*, April 2000, 67 (2), 213–234.
- Bertocchi, Graziella and Michael Spagat**, “Growth under Uncertainty with Experimentation,” *Journal of Economic Dynamics and Control*, September 1998, 23 (2), 209–231.
- Bolton, Patrick and Christopher Harris**, “Strategic Experimentation,” *Source: Econometrica Econometrica*, 1999, 67 (2), 349–374.

- Bubeck, Sébastien and Nicolò Cesa-Bianchi**, “Regret Analysis of Stochastic and Non-stochastic Multi-Armed Bandit Problems,” 2012, 5 (1), 1–122.
- Caro, Felipe and Jérémie Gallien**, “Dynamic Assortment with Demand Learning for Seasonal Consumer Goods,” *Management Science*, 2007, 53 (2), 276–292.
- Chapelle, Olivier and Lihong Li**, “An Empirical Evaluation of Thompson Sampling,” in “Advances in Neural Information Processing Systems (NIPS)” 2011.
- Chernew, Michael, Gautam Gowrisankaran, and Dennis P. Scanlon**, “Learning and the Value of Information: Evidence from Health Plan Report Cards,” *Journal of Econometrics*, May 2008, 144 (1), 156–174.
- Ching, Andrew and Masakazu Ishihara**, “The Effects of Detailing on Prescribing Decisions under Quality Uncertainty,” *Quantitative Marketing and Economics*, June 2010, 8 (2), 123–165.
- Ching, Andrew T**, “Consumer Learning and Heterogeneity: Dynamics of Demand for Prescription Drugs after Patent Expiration,” *International Journal of Industrial Organization*, 2010, 28, 619–638.
- , **Tülin Erdem, and Michael P Keane**, “Learning Models: An Assessment of Progress, Challenges, and New Developments,” *Marketing Science*, 2013, 32 (6).
- Coscelli, Andrea and Matthew Shum**, “An Empirical Model of Learning and Patient Spillovers in New Drug Entry,” *Journal of Econometrics*, October 2004, 122 (2), 213–246.
- Crawford, Gregory S. and Matthew Shum**, “Uncertainty and Learning in Pharmaceutical Demand,” *Econometrica*, 2005, 73 (4), 1137–1173.
- Creane, Anthony**, “Experimentation with Heteroskedastic Noise,” *Economic Theory*, March 1994, 4 (2), 275–286.
- Deneckere, R., H. P. Marvel, and J. Peck**, “Demand Uncertainty, Inventories, and Resale Price Maintenance,” *The Quarterly Journal of Economics*, August 1996, 111 (3), 885–913.
- Deneckere, Raymond, Howard P Marvel, and James Peck**, “Demand Uncertainty and Price Maintenance: Markdowns as Destructive Competition,” *Source: The American Economic Review*, 1997, 87 (4), 619–641.
- Dietvorst, Berkeley J, Joseph P Simmons, and Cade Massey**, “Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err,” *Journal of Experimental Psychology: General*, 2015, 144 (1), 114–126.
- , –, –, and –, “Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them,” *Management Science*, 2016.
- Dzindolet, Mary T., Linda G. Pierce, Hall P. Beck, and Lloyd A. Dawe**, “The Perceived Utility of Human and Automated Aids in a Visual Detection Task,” *Human Factors: The Journal of the Human Factors and Ergonomics Society*, March 2002, 44 (1), 79–94.
- Erdem, Tülin, Michael P Keane, and Baohong Sun**, “A Dynamic Model of Brand Choice When Price and Advertising Signal Product Quality,” *Marketing Science*, 2008, 27 (6).
- , **Michael P. Keane, T. Sabri Öncü, and Judi Strebel**, “Learning About Computers: An Analysis of Information Search and Technology Choice,” *Quantitative Marketing and Economics*, September 2005, 3 (3), 207–247.
- Fehr, Ernst, Holger Herz, and Tom Wilkening**, “The Lure of Authority: Motivation

- and Incentive Effects of Power,” *American Economic Review*, 2013, *103* (4), 1325–1359.
- Gallego, Guillermo and Huseyin Topaloglu**, “Constrained Assortment Optimization for the Nested Logit Model,” *Management Science*, 2014, *60* (10).
- Gittins, J. and D. Jones**, “A Dynamic Allocation Index for the Sequential Design of Experiments,” in J. Gani, ed., *Progress in Statistics*, Amsterdam: North-Holland, 1974, pp. 241–266.
- Harpaz, Giora, Wayne Y. Lee, and Robert L. Winkler**, “Learning, Experimentation, and the Optimal Output Decisions of a Competitive Firm,” *Management Science*, June 1982, *28* (6), 589–603.
- Hitsch, Günter J.**, “An Empirical Model of Optimal Dynamic Product Launch and Exit Under Demand Uncertainty,” *Marketing Science*, January 2006, *25* (1), 25–50.
- Hoffman, Mitchell, Lisa B Kahn, and Danielle Li**, “Discretion in Hiring,” *The Quarterly Journal of Economics*, May 2018, *133* (2), 765–800.
- Israel, Mark**, “Services as Experience Goods: An Empirical Examination of Consumer Learning in Automobile Insurance,” *American Economic Review*, November 2005, *95* (5), 1444–1463.
- Jr., Joseph E. Harrington**, “Experimentation and Learning in a Differentiated-Products Duopoly,” *Journal of Economic Theory*, June 1995, *66* (1), 275–288.
- Keller, Godfrey and Sven Rady**, “Strategic Experimentation with Poisson Bandits,” *Theoretical Economics*, 2010, *5* (2), 275–311.
- , –, **Martin Cripps, and Martin Cripps**¹, “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 2005, *73* (1), 39–68.
- May, Benedict C, Nathan Korda, Anthony Lee, and David S Leslie**, “Optimistic Bayesian Sampling in Contextual-Bandit Problems,” *Journal of Machine Learning Research*, 2012, pp. 2069–2106.
- Mirman, Leonard J., Larry Samuelson, and Amparo Urbano**, “Monopoly Experimentation,” *International Economic Review*, August 1993, *34* (3), 549.
- Morita, Tetsuro, Yuya Takano, Takashi Nakajima, Hiroshi Go, and Hiromichi Yasuoka**, “New Marketing Strategy Using the Electronic Money at Beverage Vending Machines (in Japanese),” Technical Report 2012.
- Narayanan, Sridhar and Puneet Manchanda**, “Heterogeneous Learning and the Targeting of Marketing Communication for New Products,” *Marketing Science*, 2009, *28* (3).
- , –, and **Pradeep K. Chintagunta**, “Temporal Differences in the Role of Marketing Communication in New Product Categories,” *Journal of Marketing Research*, August 2005, *42* (3), 278–290.
- Önköl, Dilek, Paul Goodwin, Mary Thomson, Sinan Gönöl, and Andrew Pollock**, “The Relative Influence of Advice from Human Experts and Statistical Methods on Forecast Adjustments,” *Journal of Behavioral Decision Making*, October 2009, *22* (4), 390–409.
- Owens, David, Zachary Grossman, and Ryan Fackler**, “The Control Premium: A Preference for Payoff Autonomy,” *American Economic Journal*, 2014, *6* (4), 138–161.
- Prahl, Andrew and Lyn Van Swol**, “Understanding Algorithm Aversion: When Is Advice from Automation Discounted?,” *Journal of Forecasting*, 2017, *36* (6), 691–702.
- Raman, Kalyan and Rabikar Chatterjee**, “Optimal Monopolist Pricing Under Demand Uncertainty in Dynamic Markets,” *Management Science*, January 1995, *41* (1), 144–162.

- Rob, Rafael**, “Learning and Capacity Expansion under Demand Uncertainty,” *The Review of Economic Studies*, June 1991, *58* (4), 655.
- Rusmevichientong, Paat, David Shmoys, Chaoxu Tong, and Huseyin Topaloglu**, “Assortment Optimization under the Multinomial Logit Model with Random Choice Parameters,” *Production and Operations Management*, 2014, *23* (11), 2023–2039.
- , **Zuo-Jun Max Shen, and David B Shmoys**, “Dynamic Assortment Optimization with a Multinomial Logit Choice Model and Capacity Constraint,” *Operations Research*, 2010, *58* (6), 1666–1680.
- Scott, Steven L**, “A Modern Bayesian Look at the Multi-Armed Bandit,” *Applied Stochastic Models in Business and Industry*, 2010, *26*, 639–658.
- Szymanowski, Maciej and Els Gijbrecchts**, “Consumption-Based Cross-Brand Learning: Are Private Labels Really Private?,” *Journal of Marketing Research*, April 2012, *49* (2), 231–246.
- Thompson, William**, “On the Likelihood That One Unknown Probability Exceeds in View of the Evidence of Two Samples,” *Biometrika*, 1933, *25* (3/4), 285–294.
- Wang, Ruxian**, “Capacitated Assortment and Price Optimization under the Multinomial Logit Model,” November 2012, *40* (6), 492–497.

Figure 2: A JRE Beverage Vending Machine

Figure 3: A JRE Vending Machine Operation Worker

Figure 4: Fit of the Demand Function

(a) In-sample fit

(b) One-period-ahead forecasting fit

¹ The unit of observation is vending machine, week, and product. The scale is the log difference from the outside share. Because the sample size is huge, I randomly sampled 10,000 observations to display the plot.

Figure 5: Mean TS and Actual Regret Over Time

¹ The dashed lines are 99% confidence intervals of the mean regrets.

B Table

Table 1: Number of Unique Values

<i>Semester</i>	<i>Product</i>	<i>Machine</i>	<i>Station</i>	<i>Area</i>	<i>Subcontractor</i>
2013S1	386	4841	592	184	10
2013S2	364	4640	538	183	8
2014S1	374	4709	536	183	8
2014S2	367	4769	537	183	8
2015S1	376	4884	538	183	9
2015S2	350	4918	539	183	9
2016S1	336	5014	539	183	9
2016S2	349	4931	538	183	9

Table 2: Summary Statistics of Sales

	<i>N</i>	<i>Mean</i>	<i>Sd</i>	<i>Min</i>	<i>Max</i>	<i>Q1</i>	<i>Q2</i>	<i>Q3</i>
Sales unit	34696744	17.924	17.507	0	608	6	13	24
Sales value	34696744	2455.320	2425.749	0	302850	840	1800	3250

The unit of observation is vending machine, product, and week.

Table 3: Weekly Number of Newly Added Products During a Season

t	N	$Mean$	Sd	Min	Max	$Q1$	$Q2$	$Q3$
2	31237	1.995	4.203	0	32	0	1	2
3	31357	2.318	4.137	0	34	0	1	2
4	31472	2.145	3.519	0	32	0	1	2
5	31566	2.215	3.521	0	37	0	1	3
6	31642	2.285	3.537	0	30	0	1	3
7	31669	2.312	3.410	0	34	0	1	3
8	31694	2.536	3.791	0	36	0	1	3
9	31810	2.619	4.016	0	38	0	1	3
10	31916	2.211	3.165	0	35	0	1	3
11	31939	1.710	3.186	0	39	0	0	2
12	31992	1.565	3.015	0	33	0	0	2
13	32025	3.562	3.646	0	38	0	3	5
14	32057	1.551	2.661	0	34	0	1	2
15	32075	1.177	2.819	0	33	0	0	1
16	32133	0.992	2.541	0	37	0	0	1
17	32158	0.876	2.479	0	35	0	0	1
18	32163	0.911	2.167	0	32	0	0	1
19	32167	0.769	2.228	0	35	0	0	1
20	32160	0.764	2.339	0	34	0	0	1
21	32189	0.585	1.979	0	35	0	0	0
22	32240	0.501	1.898	0	39	0	0	0
23	32299	0.604	2.053	0	37	0	0	0
24	32318	0.710	2.352	0	35	0	0	0
25	32279	0.620	2.200	0	33	0	0	0
26	23434	0.393	1.530	0	35	0	0	0
27	18548	0.586	1.898	0	36	0	0	1

The unit of observation is vending machine and week. The data cover the period between 2013 Summer/Spring and 2016 Summer/Spring. It only counts the products on delegated slots.

Table 4: Mean Regret under TS and Actual Product Assortments Over Time

	<i>TS</i>	<i>Actual</i>
Week 1	0.11 [0.11; 0.12]	0.15 [0.15; 0.16]
Week 2	0.05 [0.04; 0.05]	0.12 [0.11; 0.12]
Week 3	0.04 [0.03; 0.04]	0.11 [0.11; 0.11]
Week 4	0.03 [0.03; 0.03]	0.10 [0.10; 0.11]
Week 5	0.03 [0.03; 0.03]	0.10 [0.09; 0.10]
Week 6	0.03 [0.03; 0.03]	0.10 [0.09; 0.10]
Week 7	0.03 [0.03; 0.03]	0.09 [0.09; 0.10]
Week 8	0.03 [0.03; 0.03]	0.09 [0.09; 0.10]
Week 9	0.03 [0.03; 0.04]	0.09 [0.09; 0.09]
Week 10	0.03 [0.03; 0.03]	0.09 [0.08; 0.09]
Week 11	0.03 [0.03; 0.03]	0.08 [0.08; 0.09]
Week 12	0.03 [0.03; 0.03]	0.08 [0.08; 0.09]
Week 13	0.03 [0.03; 0.04]	0.09 [0.08; 0.09]
Week 14	0.03 [0.03; 0.03]	0.08 [0.08; 0.08]
Week 15	0.03 [0.03; 0.03]	0.08 [0.07; 0.08]
Week 16	0.03 [0.03; 0.03]	0.08 [0.07; 0.08]
Week 17	0.03 [0.03; 0.03]	0.08 [0.07; 0.08]
Week 18	0.03 [0.03; 0.03]	0.08 [0.07; 0.08]
Week 19	0.03 [0.03; 0.03]	0.08 [0.07; 0.08]
Week 20	0.03 [0.03; 0.03]	0.07 [0.07; 0.08]
Week 21	0.03 [0.03; 0.03]	0.07 [0.07; 0.08]
Week 22	0.03 [0.03; 0.03]	0.07 [0.07; 0.08]
Week 23	0.03 [0.03; 0.03]	0.07 [0.07; 0.08]
Week 24	0.03 [0.03; 0.03]	0.07 [0.07; 0.08]
Week 25	0.03 [0.03; 0.03]	0.07 [0.07; 0.08]
Week 26	0.03 [0.02; 0.03]	0.07 [0.06; 0.07]
Week 27	0.03 [0.03; 0.03]	0.07 [0.06; 0.07]
R ²	0.68	0.70
Adj. R ²	0.68	0.70
Num. obs.	33850	33850
RMSE	0.03	0.06

¹ 99% confidence intervals are reported.

Table 5: Prior Belief for New Products

<i>Setting</i>	<i>N</i>	<i>Mean</i>	<i>Sd</i>	<i>Min</i>	<i>Max</i>	<i>Q1</i>	<i>Q2</i>	<i>Q3</i>
October 17, 16°C	3276	0.775	0.345	0.1	5.8	0.5	0.7	1.0
November 21, 8°C	910	0.742	0.328	0.1	2.0	0.5	0.7	1.0
December 19, 5°C	2002	0.894	0.428	0.1	3.8	0.6	0.8	1.1

¹ The summary statistics is across products, and N is the number of relevant products.

Table 6: Predictive Power of Worker's Prior Belief for New Goods

	(1)	(2)	(3)
Constant	-2.61*** (0.02)	-5.87*** (0.01)	-2.56*** (0.02)
Average Belief	0.88*** (0.00)		0.86*** (0.00)
Worker Belief		0.29*** (0.00)	0.03*** (0.00)
R ²	0.35	0.10	0.35
Adj. R ²	0.35	0.10	0.35
Num. obs.	204204	204204	204204
RMSE	1.15	1.36	1.15

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are given in parentheses.

Figure 6: An Example of Instruction to a Worker

	A	B	Hot/cold	Category	D	Product name	E	Product code I	F	Product code II	G	Machine code	H	Current slots	I	Target slots	J	Centralized slots	K	Target change	L	M
1		ルート名		HC区分	大分類	商品名		JANコード		商品コード		指定		現在のコラム数		目標コラム数		うち指定コラム数		適正コラム増減数		
2	1			コールド	緑茶							C		5.3		8		4		2.8		
3	2			コールド	緑茶									0		2.4		0		2.4		
4	3			コールド	緑茶							A, B, D, E, Z, X		21		22		22		0.7		
5	4			コールド	緑茶							A		13.7		12.2		6		-1.5		
6	5			コールド	緑茶									4.2		0.6		0		-3.6		
7	6			コールド	緑茶									6.3		0		0		-6.3		
8	7																					
9	8			コールド	その他茶 (緑茶以外)							C		8.4		19.1		4		10.7		
10	9			コールド	その他茶 (緑茶以外)									0		7.1		0		7.1		
11	10			コールド	その他茶 (緑茶以外)									4.2		3.7		0		-0.5		
12	11			コールド	その他茶 (緑茶以外)							B		5.3		4		4		-1.5		
13	12			コールド	その他茶 (緑茶以外)									14.7		12.9		0		-1.8		
14	13			コールド	その他茶 (緑茶以外)									2.1		0		0		-2.1		
15	14			コールド	その他茶 (緑茶以外)									5.3		0		0		-5.3		
16	15			コールド	その他茶 (緑茶以外)							A		11.6		6		6		-6		
17	16			コールド	その他茶 (緑茶以外)									7.4		0		0		-7.4		
18	17			コールド	その他茶 (緑茶以外)									12.6		4.4		0		-8.3		
19	18			コールド	その他茶 (緑茶以外)							D, E		15.8		6.2		6		-9.6		
20	19			コールド	その他茶 (緑茶以外)							B		14.7		4		4		-11		
21	20																					
22	21			コールド	ミネラルウォーター									0		4.1		0		4.1		
23	22			コールド	ミネラルウォーター							C		1.1		4.5		4		3.4		
24	23			コールド	ミネラルウォーター									5.3		5.3		0		0.1		
25	24			コールド	ミネラルウォーター							A, B, C, D, Z		25.3		22.6		16		-2.6		
26	25			コールド	ミネラルウォーター							A, B, D, E, Z, X		25.3		22.5		22		-2.8		
27	26			コールド	ミネラルウォーター							A		8.4		6		6		-2.8		
28	27			コールド	ミネラルウォーター							D		5.3		2		2		-3.4		
29	28			コールド	ミネラルウォーター							C, X		15.8		10		10		-6.4		

¹ The red annotations are translations for column names. The area and product names are black-lacquered because they are confidential. Product code I is an universal ID and product code II is an internal ID. Machine code refers to vending machine types in which the product is required. The number of current slots can be fractional because sometimes, the product was filled only for a few days during a week. The number of target slots can be fractional because the algorithm is stochastic and the advice is based on the Monte Carlo average.

Table 7: Effect of Algorithmic Advice on Assortment Decisions

	(1)	(2)	(3)
Net Num. Targeted	0.327*** (0.006)	0.113*** (0.005)	0.118*** (0.005)
Net Num. Targeted $\times$ Advice	-0.010 (0.008)	0.001 (0.006)	0.001 (0.006)
Net Num. Targeted $\times$ Integration	0.016** (0.008)	0.014** (0.006)	0.014** (0.006)
Num. Centralized	1.092*** (0.003)	0.836*** (0.004)	0.841*** (0.004)
R ²	0.744	0.841	0.843
Adj. R ²	0.744	0.840	0.842
Num. obs.	74029	74029	74029
RMSE	5.842	4.607	4.586

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are given in parentheses. Further, Model 2 controls for product-specific fixed effects, and Model 3 controls for year-week-specific fixed effects.

Table 8: Effect of Algorithmic Advice on Normalized Sales

	(1)	(2)	(3)	(4)
Advice	-0.027 (0.051)	-0.027 (0.047)	-0.027 (0.051)	-0.027 (0.047)
Integration	0.017 (0.051)	0.017 (0.046)	0.017 (0.051)	0.017 (0.046)
R ²	0.000	0.187	0.000	0.187
Adj. R ²	-0.002	0.182	-0.002	0.182
Num. obs.	1092	1092	1092	1092
RMSE	0.694	0.627	0.694	0.627

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are given in parentheses. The dependent variable for Models (1) and (2) is sales volume normalized by last year's average sales volume in the area. The dependent variable for Models (3) and (4) is sales value normalized by last year's average sales value in the area. Models (2) and (4) control for month-specific fixed effects.

Table 9: Effect of Algorithmic Advice on Assortment Decisions: Area/Worker- and Machine-level Heterogeneity

	(1)	(2)	(3)	(4)	(5)
Net Num. Targeted	0.117*** (0.005)	0.112*** (0.005)	0.106*** (0.008)	0.109*** (0.008)	0.023*** (0.006)
Net Num. Targeted $\times$ Regret	-0.006 (0.010)	-0.008 (0.010)	-0.003 (0.010)	0.001 (0.010)	-0.002 (0.006)
Net Num. Targeted $\times$ Delegation				-0.013** (0.006)	-0.018*** (0.003)
Net Num. Targeted $\times$ Experience			0.008 (0.010)	0.012 (0.010)	0.045*** (0.006)
Net Num. Targeted $\times$ Area Sales Volatility		0.021*** (0.005)	0.020*** (0.005)	0.018*** (0.005)	0.009*** (0.003)
Net Num. Targeted $\times$ High Sales Volume Machine					-0.006 (0.007)
Net Num. Targeted $\times$ High Sales Volatility Machine					-0.033*** (0.006)
Net Num. Targeted $\times$ Advice	0.004 (0.006)	0.008 (0.006)	0.008 (0.010)	0.022** (0.010)	0.080*** (0.008)
Net Num. Targeted $\times$ Advice $\times$ Regret	0.046*** (0.013)	0.047*** (0.013)	0.043*** (0.013)	0.026** (0.013)	0.019** (0.008)
Net Num. Targeted $\times$ Advice $\times$ Delegation				-0.029*** (0.008)	-0.008* (0.005)
Net Num. Targeted $\times$ Advice $\times$ Experience			0.001 (0.014)	-0.018 (0.014)	-0.048*** (0.009)
Net Num. Targeted $\times$ Advice $\times$ Area Sales Volatility		-0.014** (0.007)	-0.012* (0.007)	-0.019** (0.008)	-0.032*** (0.005)
Net Num. Targeted $\times$ Advice $\times$ High Sales Volume Machine					-0.083*** (0.010)
Net Num. Targeted $\times$ Advice $\times$ High Sales Volatility Machine					0.045*** (0.010)
Net Num. Targeted $\times$ Integration	0.010 (0.006)	0.001 (0.006)	-0.018* (0.009)	-0.039*** (0.010)	-0.068*** (0.008)
Net Num. Targeted $\times$ Integration $\times$ Regret	-0.064*** (0.011)	-0.079*** (0.011)	-0.073*** (0.011)	-0.063*** (0.011)	-0.036*** (0.007)
Net Num. Targeted $\times$ Integration $\times$ Delegation				0.053*** (0.007)	0.010** (0.005)
Net Num. Targeted $\times$ Integration $\times$ Experience			0.032** (0.013)	0.047*** (0.013)	0.021** (0.008)
Net Num. Targeted $\times$ Integration $\times$ Area Sales Volatility		0.043*** (0.006)	0.045*** (0.007)	0.055*** (0.007)	0.036*** (0.004)
Net Num. Targeted $\times$ Integration $\times$ High Sales Volume Machine					0.083*** (0.010)
Net Num. Targeted $\times$ Integration $\times$ High Sales Volatility Machine					-0.019** (0.009)
Num. Centralized	0.841*** (0.004)	0.840*** (0.004)	0.837*** (0.004)	0.837*** (0.004)	0.246*** (0.002)
R ²	0.843	0.843	0.844	0.844	0.566
Adj. R ²	0.842	0.842	0.843	0.844	0.565
Num. obs.	74029	74029	72753	72753	145688
RMSE	4.585	4.580	4.586	4.584	4.179

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are given in parentheses. Standard errors are clustered at the area level. Product-specific fixed effects and year-week-specific fixed effects are controlled.

Table 10: Effect of Algorithmic Advice on Normalized Sales Volume: Area/Worker-level Heterogeneity

	(1)	(2)	(3)	(4)	(5)
Advice $\times$ Regret	0.100*** (0.033)	0.030** (0.015)	0.026* (0.015)	0.003 (0.014)	-0.044 (0.048)
Advice $\times$ Delegation				-0.018* (0.011)	0.002 (0.035)
Advice $\times$ Experience			-0.035 (0.024)	-0.015 (0.021)	-0.097 (0.071)
Advice $\times$ Area Sales Volatility		-0.005 (0.012)	-0.001 (0.013)	-0.000 (0.011)	0.032 (0.038)
Advice $\times$ High Sales Volume Machine					-0.184*** (0.071)
Advice $\times$ High Sales Volatioity Machine					0.046 (0.071)
Integration	0.008 (0.025)	0.005 (0.011)	-0.047*** (0.016)	-0.043*** (0.014)	-0.141** (0.062)
Integration $\times$ Regret	-0.024 (0.040)	-0.034* (0.018)	-0.011 (0.018)	-0.024 (0.016)	-0.056 (0.055)
Integration $\times$ Delegation				0.044*** (0.011)	0.011 (0.036)
Integration $\times$ Experience			0.108*** (0.023)	0.101*** (0.021)	0.033 (0.069)
Integration $\times$ Area Sales Volatility		-0.002 (0.011)	0.007 (0.012)	0.004 (0.010)	-0.023 (0.037)
Integration $\times$ High Sales Volume Machine					0.182** (0.071)
Integration $\times$ High Sales Volatility Machine					-0.037 (0.070)
R ²	0.854	0.970	0.972	0.978	0.647
Adj. R ²	0.853	0.970	0.971	0.977	0.644
Num. obs.	1092	1092	1068	1068	3060
RMSE	0.335	0.152	0.149	0.132	0.734

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are given in parentheses. The dependent variables are the sales volume normalized by last year's average sales volume in the unit. The month-specific fixed effects are controlled.

Table 11: How Do Actual and TS Product Assortments Differ?

<i>Product Type</i>	<i>Control</i>	<i>Advice</i>	<i>Integration</i>
Constant	-0.232*** (0.004)		
New	0.081*** (0.008)	-0.025** (0.010)	-0.011 (0.010)
Private Brand	0.274*** (0.005)	0.024*** (0.006)	0.028*** (0.006)
Hot	0.292*** (0.006)	-0.004 (0.008)	-0.111*** (0.008)
Aluminium Bottle	0.480*** (0.012)	-0.080*** (0.017)	-0.038** (0.017)
Aluminium Can	-0.182*** (0.014)	0.097*** (0.020)	0.031 (0.020)
Volume	0.032*** (0.003)	-0.010** (0.004)	0.006 (0.004)
CocaByCoca	-0.100*** (0.009)	0.065*** (0.013)	0.001 (0.013)
Coca	0.218*** (0.008)	-0.066*** (0.010)	0.012 (0.010)
Price	-0.049*** (0.003)	0.046*** (0.004)	-0.005 (0.004)

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are given in parentheses. The dependent variables are the difference in the actual and TS product assortment shares in percentage.

Table 12: Regret During the Experiment Across Areas

	<i>N</i>	<i>Mean</i>	<i>Sd</i>	<i>Min</i>	<i>Max</i>	<i>Q1</i>	<i>Q2</i>	<i>Q3</i>
Actual regret (%)	18128	0.120	0.044	0.019	0.374	0.089	0.113	0.143
TS regret (%)	18128	0.053	0.031	0.000	0.258	0.031	0.048	0.070
Potential change (% points)	18128	0.068	0.051	-0.059	0.363	0.035	0.060	0.093

The unit of observation is areas during the experiment period.

Table 13: What Needs To Be Improved For Workers to Follow an Algorithm? (Multiple Answers)

(a) Unweighted				
<i>Answer</i>	<i>All</i>	<i>Control</i>	<i>Treatment I</i>	<i>Treatment II</i>
a. Reliable demand prediction	76	22	22	32
b. Integration of workers' prediction	52	16	21	15
c. Integration of managers' prediction	5	2	1	2
d. Detailed explanation of the algorithm	54	19	16	19
e. Coordination with manager's order	18	8	5	5
f. Reduced burden of the work	101	33	31	37
g. Machine-level instruction	58	16	22	20
h. Inside-station-level instruction	23	6	10	7
i. Station-level instruction	43	12	12	19
j. Mobile interface	33	12	9	12
k. Visual interface	45	15	14	16
l. Coordination with product abolition schedule	74	26	21	27
m. Compensation to the loss	15	7	1	7

(b) Weighted				
<i>Answer</i>	<i>All</i>	<i>Control</i>	<i>Treatment I</i>	<i>Treatment II</i>
a. Reliable demand prediction	21.1	6.0	6.4	8.6
b. Integration of workers' prediction	15.4	4.5	6.6	4.3
c. Integrateion of managers' prediction	1.1	0.8	0.1	0.2
d. Detailed explanation of the algorithm	15.0	5.4	4.9	4.7
e. Coordination with manager's order	5.7	2.5	2.0	1.2
f. Reduced burden of the work	32.2	10.9	10.0	11.2
g. Machine-level instruction	17.1	4.8	7.5	4.8
h. Inside-station-level instruction	3.8	0.9	1.8	1.1
i. Station-level instruction	11.2	3.5	3.1	4.6
j. Mobile interface	8.8	2.7	3.1	3.1
k. Visual interface	10.4	3.3	3.7	3.4
l. Coordination with product abolition schedule	19.2	7.3	5.7	6.3
m. Compensation to the loss	3.0	1.5	0.1	1.4

¹ The unit of observations is an area. Multiple answers are allowed. The unweighted number sums up the number of areas whose answer includes the item. For the weighted number, if the answer of an area includes the item, then divide by the number of items included in the answer of the area, and then, sum up the weighted numbers across areas.

C Online Appendix

Figure A1: Weekly Trend in the Share of Hot Product Slots Per Machine

¹ A point is the average share of hot product slots per machine. The data cover the period between 2013 Summer/Spring and 2016 Summer/Spring.

Figure A2: Mean TS and Actual Success Rates Over Time

¹ The dashed lines indicate 99% confidence intervals of the mean success rates.

Table A1: Summary Statistics of Product Assortments Per Machine

	<i>N</i>	<i>Mean</i>	<i>Sd</i>	<i>Min</i>	<i>Max</i>	<i>Q1</i>	<i>Q2</i>	<i>Q3</i>
Num. slots	869664	35.369	4.045	5	42.000	34.000	36.000	36.000
Num. products	869664	32.298	4.410	1	42.000	30.000	33.000	35.000
Num. categories	869664	11.073	0.917	1	12.000	11.000	11.000	12.000
Num. hot product	869664	8.335	8.389	0	32.000	0.000	6.000	18.000
Share hot product	869664	0.237	0.239	0	1.000	0.000	0.167	0.500
Num. new product	869664	7.393	4.489	0	24.000	4.000	8.000	11.000
Share new product	869664	0.209	0.125	0	0.676	0.111	0.222	0.306
Num. private product	869664	0.238	0.689	0	9.000	0.000	0.000	0.000
Share private product	869664	0.007	0.020	0	0.250	0.000	0.000	0.000
Num. delegated	869664	17.689	6.996	0	42.000	13.000	17.000	22.000
Share delegated	869664	0.498	0.187	0	1.000	0.361	0.467	0.611

The unit of observation is vending machine and week. The data cover the period between 2013 Summer/Spring and 2016 Summer/Spring.

Table A2: Summary Statistics of Product Sales Per Machine

	<i>N</i>	<i>Mean</i>	<i>Sd</i>	<i>Min</i>	<i>Max</i>	<i>Q1</i>	<i>Q2</i>	<i>Q3</i>
Num. hot product	869664	147.519	217.017	0	4247.0	0.000	55.000	225.000
Share hot product	869664	0.273	0.290	0	1.0	0.000	0.165	0.563
Num. new product	869664	127.757	140.532	0	3341.0	32.000	85.000	176.000
Share new product	869664	0.202	0.132	0	1.0	0.099	0.195	0.290
Num. private product	869664	3.979	15.315	0	510.0	0.000	0.000	0.000
Share private product	869664	0.006	0.019	0	0.4	0.000	0.000	0.000
Num. delegated	869664	288.749	275.122	0	7978.0	113.000	211.000	376.000
Share delegated	869664	0.467	0.199	0	1.0	0.323	0.426	0.586

The unit of observation is vending machine and week. The data cover the period between 2013 Summer/Spring and 2016 Summer/Spring.

Figure A3: Residual Plots of Demand Estimation

¹ The unit of observation is vending machine, week, and product. The scale is the log difference from the outside share. Because the sample size is huge, I randomly sampled 10,000 observations to display the plot.

Table A3: Summary Statistics of Product Assortments Per Machine By Category

	<i>N</i>	<i>Mean</i>	<i>Sd</i>	<i>Min</i>	<i>Max</i>	<i>Q1</i>	<i>Q2</i>	<i>Q3</i>
Num. Green tea	869664	2.758	0.975	0	26.000	2.000	3.000	3.000
Share Green tea	869664	0.078	0.027	0	0.743	0.056	0.083	0.088
Num. Other tea	869664	4.421	1.521	0	17.0	3.0	4.000	5.000
Share Other tea	869664	0.125	0.040	0	0.5	0.1	0.119	0.143
Num. Water	869664	3.246	1.146	0	36	2.000	3.000	4.000
Share Water	869664	0.092	0.032	0	1	0.067	0.086	0.111
Num. Tonic water	869664	2.494	1.344	0	12.0	1.000	2.000	3.000
Share Tonic water	869664	0.070	0.037	0	0.4	0.048	0.067	0.095
Num. Canned coffee	869664	5.776	2.031	0	36	4.000	6.000	7.000
Share Canned coffee	869664	0.164	0.056	0	1	0.133	0.167	0.194
Num. Bottled coffee	869664	1.803	1.06	0	10.000	1.000	2.000	2.000
Share Bottled coffee	869664	0.051	0.03	0	0.278	0.028	0.048	0.067
Num. English tea	869664	2.058	1.213	0	10.0	1.000	2.000	3.000
Share English tea	869664	0.058	0.034	0	0.3	0.029	0.056	0.083
Num. Carbonated water	869664	4.954	2.421	0	36	3.000	5.000	6.000
Share Carbonated water	869664	0.140	0.065	0	1	0.095	0.139	0.167
Num. Fruits	869664	3.279	1.387	0	18.000	2.000	3.000	4.000
Share Fruits	869664	0.092	0.037	0	0.429	0.067	0.088	0.111
Num. Healthy drink	869664	1.702	0.786	0	30	1.000	2.000	2.000
Share Healthy drink	869664	0.048	0.022	0	1	0.029	0.056	0.056
Num. Specialty	869664	0.594	0.766	0	6.000	0	0	1.000
Share Specialty	869664	0.017	0.022	0	0.167	0	0	0.028
Num. Other	869664	2.283	1.563	0	28	1.000	2.000	3.000
Share Other	869664	0.065	0.044	0	1	0.028	0.057	0.095

The unit of observation is vending machine and week. The data cover the period between 2013 Summer/Spring and 2016 Summer/Spring.

Table A4: Summary Statistics of Product Sales Per Machine By Category

	<i>N</i>	<i>Mean</i>	<i>Sd</i>	<i>Min</i>	<i>Max</i>	<i>Q1</i>	<i>Q2</i>	<i>Q3</i>
Num. Green tea	869664	51.469	59.462	0	2712	19.000	37.000	66.000
Share Green tea	869664	0.080	0.035	0	1	0.057	0.075	0.097
Num. Other tea	869664	81.468	89.527	0	3720	27.000	56.000	106.000
Share Other tea	869664	0.119	0.045	0	1	0.089	0.117	0.147
Num. Water	869664	74.135	79.123	0	2112	23.000	50.000	97.000
Share Water	869664	0.108	0.047	0	1	0.073	0.105	0.138
Num. Tonic water	869664	46.904	59.572	0	1418	11.000	27.000	59.000
Share Tonic water	869664	0.067	0.043	0	1	0.034	0.062	0.095
Num. Canned coffee	869664	89.147	79.838	0	2514	41.000	71.000	115.000
Share Canned coffee	869664	0.158	0.073	0	1	0.108	0.146	0.194
Num. Bottled coffee	869664	36.485	35.778	0	644	13.000	27.000	49.000
Share Bottled coffee	869664	0.061	0.038	0	1	0.034	0.055	0.083
Num. English tea	869664	34.428	36.645	0	617.00	10.000	23.00	46.000
Share English tea	869664	0.057	0.040	0	0.52	0.027	0.05	0.081
Num. Carbonated water	869664	87.438	92.299	0	2234	29.000	59.00	113.000
Share Carbonated water	869664	0.136	0.075	0	1	0.084	0.13	0.179
Num. Fruits	869664	53.700	55.838	0	1053	18.000	37.000	71.000
Share Fruits	869664	0.081	0.038	0	1	0.055	0.079	0.104
Num. Healthy drink	869664	28.718	30.415	0	524	8.000	19.000	39.000
Share Healthy drink	869664	0.045	0.029	0	1	0.023	0.041	0.062
Num. Specialty	869664	9.203	18.877	0	446.000	0	0	11.000
Share Specialty	869664	0.014	0.022	0	0.362	0	0	0.023
Num. Other	869664	41.039	45.337	0	964	10.000	28.000	57.000
Share Other	869664	0.074	0.062	0	1	0.025	0.064	0.111

The unit of observation is vending machine and week. The data cover the period between 2013 Summer/Spring and 2016 Summer/Spring.

Table A5: *Ex Post* Belief About γ and σ^2

<i>Parameter</i>	<i>N</i>	<i>Mean</i>	<i>Sd</i>	<i>Min</i>	<i>Max</i>	<i>Q1</i>	<i>Q2</i>	<i>Q3</i>
γ	1472	0.751	0.089	0.354	1.021	0.703	0.756	0.809
σ^2	1472	0.271	0.119	0.086	1.010	0.185	0.234	0.325

¹ The summary statistics is across areas and seasons, and N is the number of relevant areas and seasons.

Table A6: *Ex Post* Belief About β

<i>Group</i>	<i>N</i>	<i>Mean</i>	<i>Sd</i>	<i>Min</i>	<i>Max</i>	<i>Q1</i>	<i>Q2</i>	<i>Q3</i>
<i>Cold</i>								
Green tea	1472	0.050	0.022	0.000	0.695	0.043	0.051	0.059
Other tea	1471	0.069	0.018	-0.132	0.121	0.062	0.071	0.079
Water	1472	0.046	0.015	0.006	0.408	0.041	0.047	0.053
Tonic water	1472	0.063	0.019	0.005	0.411	0.054	0.065	0.074
Canned coffee	1470	0.079	0.034	-0.010	0.201	0.051	0.073	0.107
Bottled coffee	1471	0.060	0.034	-0.633	0.165	0.039	0.057	0.081
English tea	1471	0.048	0.035	-0.270	0.378	0.024	0.045	0.074
Carbonated water	1472	0.050	0.014	0.001	0.250	0.046	0.052	0.058
Fruit	1472	0.054	0.032	-0.317	0.978	0.041	0.055	0.067
Healthy drink	1472	0.033	0.021	-0.043	0.703	0.027	0.033	0.040
Specialty	1188	0.038	0.068	-0.303	1.275	0.023	0.046	0.060
Other	1467	0.057	0.063	-0.897	0.694	0.024	0.045	0.081
<i>Hot</i>								
Green tea	1472	-0.034	0.029	-0.209	0.462	-0.047	-0.032	-0.020
Other tea	1472	-0.055	0.039	-0.395	0.184	-0.068	-0.051	-0.035
Canned coffee	1472	-0.049	0.028	-0.354	0.091	-0.069	-0.047	-0.031
Bottled coffee	1471	-0.043	0.038	-0.403	0.260	-0.056	-0.038	-0.022
English tea	1472	-0.053	0.036	-0.336	0.143	-0.075	-0.049	-0.032
Fruit	1472	-0.055	0.042	-0.243	0.613	-0.075	-0.051	-0.033
Other	1472	-0.049	0.045	-0.260	0.611	-0.078	-0.043	-0.019

¹ The summary statistics is across areas and seasons, and N is the number of relevant areas and seasons.

Table A7: *Ex Post* Belief About ξ

<i>Product</i>	<i>N</i>	<i>Mean</i>	<i>Sd</i>	<i>Min</i>	<i>Max</i>	<i>Q1</i>	<i>Q2</i>	<i>Q3</i>
100 percentile product	125	-3.952	0.630	-5.946	0.376	-4.220	-3.957	-3.689
90 percentile product	358	-4.390	0.512	-6.108	-2.931	-4.653	-4.332	-4.063
80 percentile product	768	-4.670	0.525	-6.066	0.630	-4.956	-4.710	-4.485
70 percentile product	250	-5.456	0.385	-7.142	-4.355	-5.704	-5.452	-5.212
60 percentile product	1141	-5.663	0.381	-9.561	-4.323	-5.894	-5.652	-5.410
50 percentile product	351	-6.015	0.440	-7.683	-4.955	-6.259	-5.963	-5.734
40 percentile product	154	-6.281	0.505	-7.916	-4.927	-6.613	-6.308	-5.905
30 percentile product	259	-6.407	1.109	-11.968	-4.272	-6.730	-6.208	-5.801
20 percentile product	1274	-6.575	0.775	-9.089	-4.319	-7.196	-6.547	-5.986
10 percentile product	275	-6.788	0.833	-11.565	-5.334	-7.042	-6.618	-6.291

¹ The summary statistics is across areas and seasons, and N is the number of relevant areas and seasons. The percentile is among products that are estimated more than 100 areas and seasons. I cannot provide the product names because of confidentiality reasons.

Table A8: Mean TS and Actual Success Rates Over Time

	<i>TS</i>	<i>Actual</i>
Week 1	0.37 [0.36; 0.37]	0.29 [0.29; 0.30]
Week 2	0.39 [0.39; 0.40]	0.30 [0.29; 0.30]
Week 3	0.39 [0.38; 0.39]	0.29 [0.29; 0.30]
Week 4	0.37 [0.37; 0.38]	0.29 [0.28; 0.29]
Week 5	0.36 [0.36; 0.37]	0.28 [0.27; 0.28]
Week 6	0.35 [0.34; 0.35]	0.27 [0.26; 0.27]
Week 7	0.33 [0.33; 0.34]	0.25 [0.25; 0.26]
Week 8	0.32 [0.32; 0.33]	0.24 [0.24; 0.25]
Week 9	0.32 [0.31; 0.33]	0.23 [0.23; 0.24]
Week 10	0.32 [0.31; 0.32]	0.23 [0.23; 0.24]
Week 11	0.32 [0.31; 0.33]	0.23 [0.23; 0.24]
Week 12	0.32 [0.32; 0.33]	0.23 [0.23; 0.24]
Week 13	0.33 [0.32; 0.33]	0.25 [0.24; 0.25]
Week 14	0.31 [0.31; 0.32]	0.24 [0.24; 0.25]
Week 15	0.32 [0.31; 0.32]	0.24 [0.23; 0.24]
Week 16	0.32 [0.31; 0.32]	0.24 [0.23; 0.24]
Week 17	0.32 [0.31; 0.32]	0.23 [0.23; 0.24]
Week 18	0.32 [0.31; 0.32]	0.23 [0.22; 0.23]
Week 19	0.32 [0.31; 0.32]	0.23 [0.22; 0.23]
Week 20	0.32 [0.31; 0.32]	0.22 [0.22; 0.23]
Week 21	0.32 [0.31; 0.32]	0.22 [0.22; 0.23]
Week 22	0.32 [0.31; 0.32]	0.22 [0.22; 0.23]
Week 23	0.32 [0.31; 0.32]	0.22 [0.22; 0.23]
Week 24	0.32 [0.31; 0.32]	0.22 [0.22; 0.23]
Week 25	0.32 [0.31; 0.32]	0.22 [0.22; 0.23]
Week 26	0.31 [0.30; 0.31]	0.22 [0.22; 0.23]
Week 27	0.31 [0.30; 0.32]	0.22 [0.22; 0.23]
R ²	0.93	0.92
Adj. R ²	0.93	0.92
Num. obs.	33850	33850
RMSE	0.09	0.07

¹ 99% confidence intervals are reported.

Table A9: Effect of Algorithmic Advice on Assortment Decisions: Area/Worker-level Heterogeneity

	(1)	(2)	(3)
Net Num. Targeted	0.326*** (0.006)	0.112*** (0.005)	0.117*** (0.005)
Net Num. Targeted $\times$ Regret	-0.013 (0.012)	-0.005 (0.010)	-0.006 (0.010)
Net Num. Targeted $\times$ Advice	-0.006 (0.008)	0.004 (0.006)	0.004 (0.006)
Net Num. Targeted $\times$ Advice $\times$ Regret	0.055*** (0.016)	0.046*** (0.013)	0.046*** (0.013)
Net Num. Targeted $\times$ Integration	0.013 (0.008)	0.010* (0.006)	0.010 (0.006)
Net Num. Targeted $\times$ Integration $\times$ Regret	-0.063*** (0.014)	-0.064*** (0.011)	-0.064*** (0.011)
Num. Centralized	1.092*** (0.003)	0.836*** (0.004)	0.841*** (0.004)
R ²	0.744	0.841	0.843
Adj. R ²	0.744	0.841	0.842
Num. obs.	74029	74029	74029
RMSE	5.841	4.606	4.585

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are in parentheses. Moreover, Model 2 controls for product-specific fixed effects, and Model 3 controls for year-week-specific fixed effects. $Regret_i$ is the average of actual regret minus TS regret.

Table A10: Effect of Algorithmic Advice on Assortment Decisions: Area/Worker-level Heterogeneity

	(1)	(2)	(3)
Net Num. Targeted	0.328*** (0.006)	0.107*** (0.005)	0.112*** (0.005)
Net Num. Targeted $\times$ Regret	-0.011 (0.012)	-0.008 (0.010)	-0.008 (0.010)
Net Num. Targeted $\times$ Area Sales Volatility	-0.011* (0.006)	0.022*** (0.005)	0.021*** (0.005)
Net Num. Targeted $\times$ Advice	-0.008 (0.008)	0.008 (0.006)	0.008 (0.006)
Net Num. Targeted $\times$ Advice $\times$ Regret	0.053*** (0.016)	0.047*** (0.013)	0.047*** (0.013)
Net Num. Targeted $\times$ Advice $\times$ Area Sales Volatility	0.012 (0.009)	-0.015** (0.007)	-0.014** (0.007)
Net Num. Targeted $\times$ Integration	0.007 (0.008)	0.001 (0.006)	0.001 (0.006)
Net Num. Targeted $\times$ Integration $\times$ Regret	-0.072*** (0.014)	-0.079*** (0.011)	-0.079*** (0.011)
Net Num. Targeted $\times$ Integration $\times$ Area Sales Volatility	0.027*** (0.008)	0.043*** (0.006)	0.043*** (0.006)
Num. Centralized	1.092*** (0.003)	0.834*** (0.004)	0.840*** (0.004)
R ²	0.744	0.842	0.843
Adj. R ²	0.744	0.841	0.842
Num. obs.	74029	74029	74029
RMSE	5.840	4.601	4.580

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are given in parentheses. Model 2 controls for product-specific fixed effects, and Model 3 controls for year-week-specific fixed effects.

Table A11: Effect of Algorithmic Advice on Assortment Decisions: Area/Worker-level Heterogeneity

	(1)	(2)	(3)
Net Num. Targeted	0.331*** (0.010)	0.100*** (0.008)	0.106*** (0.008)
Net Num. Targeted $\times$ Regret	-0.010 (0.013)	-0.002 (0.010)	-0.003 (0.010)
Net Num. Targeted $\times$ Experience	-0.007 (0.013)	0.009 (0.010)	0.008 (0.010)
Net Num. Targeted $\times$ Area Sales Volatility	-0.011 (0.007)	0.021*** (0.005)	0.020*** (0.005)
Net Num. Targeted $\times$ Advice	-0.012 (0.013)	0.008 (0.010)	0.008 (0.010)
Net Num. Targeted $\times$ Advice $\times$ Regret	0.052*** (0.017)	0.043*** (0.013)	0.043*** (0.013)
Net Num. Targeted $\times$ Advice $\times$ Experience	0.009 (0.018)	-0.000 (0.014)	0.001 (0.014)
Net Num. Targeted $\times$ Advice $\times$ Area Sales Volatility	0.013 (0.009)	-0.013* (0.007)	-0.012* (0.007)
Net Num. Targeted $\times$ Integration	0.004 (0.012)	-0.018* (0.009)	-0.018* (0.009)
Net Num. Targeted $\times$ Integration $\times$ Regret	-0.072*** (0.014)	-0.073*** (0.011)	-0.073*** (0.011)
Net Num. Targeted $\times$ Integration $\times$ Experience	0.004 (0.017)	0.033** (0.013)	0.032** (0.013)
Net Num. Targeted $\times$ Integration $\times$ Area Sales Volatility	0.027*** (0.008)	0.045*** (0.007)	0.045*** (0.007)
Num. Centralized	1.093*** (0.003)	0.831*** (0.004)	0.837*** (0.004)
R ²	0.745	0.843	0.844
Adj. R ²	0.744	0.842	0.843
Num. obs.	72753	72753	72753
RMSE	5.858	4.607	4.586

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are given in parentheses. Model 2 controls for product-specific fixed effects, and Model 3 controls for year-week-specific fixed effects.

Table A12: Effect of Algorithmic Advice on Assortment Decisions: Area/Worker-level Heterogeneity

	(1)	(2)	(3)
Net Num. Targeted	0.340*** (0.010)	0.103*** (0.008)	0.109*** (0.008)
Net Num. Targeted $\times$ Regret	0.004 (0.013)	0.001 (0.010)	0.001 (0.010)
Net Num. Targeted $\times$ Delegation	-0.046*** (0.007)	-0.012** (0.006)	-0.013** (0.006)
Net Num. Targeted $\times$ Experience	0.009 (0.013)	0.013 (0.010)	0.012 (0.010)
Net Num. Targeted $\times$ Area Sales Volatility	-0.019*** (0.007)	0.019*** (0.005)	0.018*** (0.005)
Net Num. Targeted $\times$ Advice	-0.001 (0.013)	0.023** (0.010)	0.022** (0.010)
Net Num. Targeted $\times$ Advice $\times$ Regret	0.022 (0.017)	0.026** (0.013)	0.026** (0.013)
Net Num. Targeted $\times$ Advice $\times$ Delegation	-0.004 (0.010)	-0.029*** (0.008)	-0.029*** (0.008)
Net Num. Targeted $\times$ Advice $\times$ Experience	-0.024 (0.018)	-0.018 (0.014)	-0.018 (0.014)
Net Num. Targeted $\times$ Advice $\times$ Area Sales Volatility	0.010 (0.010)	-0.019*** (0.008)	-0.019** (0.008)
Net Num. Targeted $\times$ Integration	-0.003 (0.013)	-0.040*** (0.010)	-0.039*** (0.010)
Net Num. Targeted $\times$ Integration $\times$ Regret	-0.050*** (0.014)	-0.063*** (0.011)	-0.063*** (0.011)
Net Num. Targeted $\times$ Integration $\times$ Delegation	0.019** (0.009)	0.054*** (0.007)	0.053*** (0.007)
Net Num. Targeted $\times$ Integration $\times$ Experience	0.021 (0.017)	0.048*** (0.013)	0.047*** (0.013)
Net Num. Targeted $\times$ Integration $\times$ Area Sales Volatility	0.032*** (0.009)	0.056*** (0.007)	0.055*** (0.007)
Num. Centralized	1.093*** (0.003)	0.831*** (0.004)	0.837*** (0.004)
R ²	0.745	0.843	0.844
Adj. R ²	0.745	0.842	0.844
Num. obs.	72753	72753	72753
RMSE	5.854	4.604	4.584

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are given in parentheses. Model 2 controls for product-specific fixed effects, and Model 3 controls for year-week-specific fixed effects.

Table A13: Effect of Algorithmic Advice on the Assortment Decisions: Area/Worker- and Machine-level Heterogeneity

	(1)	(2)	(3)
Net Num. Targeted	0.157*** (0.006)	0.020*** (0.006)	0.023*** (0.006)
Net Num. Targeted $\times$ Regret	0.002 (0.007)	-0.002 (0.006)	-0.002 (0.006)
Net Num. Targeted $\times$ Delegation	-0.037*** (0.004)	-0.018*** (0.003)	-0.018*** (0.003)
Net Num. Targeted $\times$ Experience	0.031*** (0.007)	0.045*** (0.006)	0.045*** (0.006)
Net Num. Targeted $\times$ Area Sales Volatility	-0.019*** (0.004)	0.010*** (0.003)	0.009*** (0.003)
Net Num. Targeted $\times$ High Sales Volume Machine	0.010 (0.007)	-0.007 (0.007)	-0.006 (0.007)
Net Num. Targeted $\times$ High Sales Volatility Machine	-0.033*** (0.007)	-0.033*** (0.006)	-0.033*** (0.006)
Net Num. Targeted $\times$ Advice	0.062*** (0.008)	0.080*** (0.008)	0.080*** (0.008)
Net Num. Targeted $\times$ Advice $\times$ Regret	0.013 (0.009)	0.019** (0.008)	0.019** (0.008)
Net Num. Targeted $\times$ Advice $\times$ Delegation	0.005 (0.005)	-0.008* (0.005)	-0.008* (0.005)
Net Num. Targeted $\times$ Advice $\times$ Experience	-0.039*** (0.010)	-0.048*** (0.009)	-0.048*** (0.009)
Net Num. Targeted $\times$ Advice $\times$ Area Sales Volatility	-0.017*** (0.005)	-0.032*** (0.005)	-0.032*** (0.005)
Net Num. Targeted $\times$ Advice $\times$ High Sales Volume Machine	-0.093*** (0.011)	-0.083*** (0.010)	-0.083*** (0.010)
Net Num. Targeted $\times$ Advice $\times$ High Sales Volatility Machine	0.058*** (0.010)	0.044*** (0.010)	0.045*** (0.010)
Net Num. Targeted $\times$ Integration	-0.046*** (0.008)	-0.068*** (0.008)	-0.068*** (0.008)
Net Num. Targeted $\times$ Integration $\times$ Regret	-0.024*** (0.008)	-0.036*** (0.007)	-0.036*** (0.007)
Net Num. Targeted $\times$ Integration $\times$ Delegation	-0.011** (0.005)	0.010** (0.005)	0.010** (0.005)
Net Num. Targeted $\times$ Integration $\times$ Experience	0.006 (0.009)	0.021** (0.008)	0.021** (0.008)
Net Num. Targeted $\times$ Integration $\times$ Area Sales Volatility	0.025*** (0.005)	0.036*** (0.004)	0.036*** (0.004)
Net Num. Targeted $\times$ Integration $\times$ High Sales Volume Machine	0.099*** (0.011)	0.082*** (0.010)	0.083*** (0.010)
Net Num. Targeted $\times$ Integration $\times$ High Sales Volatility Machine	-0.037*** (0.010)	-0.019** (0.009)	-0.019** (0.009)
Num. Centralized	0.403*** (0.001)	0.243*** (0.002)	0.246*** (0.002)
R ²	0.476	0.565	0.566
Adj. R ²	0.476	0.564	0.565
Num. obs.	145688	145688	145688
RMSE	4.589	4.184	4.179

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are given in parentheses. Standard errors are clustered at the area level. Model 2 controls for product-specific fixed effects, and Model 3 controls for year-week-specific fixed effects.

Table A14: Effect of Algorithmic Advice on Assortment Decisions: Area/Worker- and Machine-level Heterogeneity and Potential Office-level Information-Sharing

	(1)	(2)	(3)
Net Num. Targeted	0.087*** (0.007)	0.022*** (0.006)	0.024*** (0.006)
Average Net Num. Targeted in the Office	0.150*** (0.006)	-0.006 (0.006)	-0.003 (0.006)
Net Num. Targeted $\times$ Regret	-0.006 (0.007)	-0.001 (0.006)	-0.002 (0.006)
Net Num. Targeted $\times$ Delegation	-0.032*** (0.004)	-0.018*** (0.003)	-0.019*** (0.003)
Net Num. Targeted $\times$ Experience	0.040*** (0.007)	0.045*** (0.006)	0.045*** (0.006)
Net Num. Targeted $\times$ Area Sales Volatility	-0.017*** (0.004)	0.010*** (0.003)	0.009*** (0.003)
Net Num. Targeted $\times$ High Sales Volume Machine	0.008 (0.007)	-0.007 (0.007)	-0.006 (0.007)
Net Num. Targeted $\times$ High Sales Volatility Machine	-0.035*** (0.007)	-0.033*** (0.006)	-0.033*** (0.006)
Net Num. Targeted $\times$ Advice	0.186*** (0.010)	0.075*** (0.009)	0.077*** (0.009)
Net Num. Targeted $\times$ Advice $\times$ Regret	0.017* (0.009)	0.019** (0.008)	0.019** (0.008)
Net Num. Targeted $\times$ Advice $\times$ Delegation	-0.013** (0.005)	-0.007 (0.005)	-0.007 (0.005)
Net Num. Targeted $\times$ Advice $\times$ Experience	-0.046*** (0.010)	-0.048*** (0.009)	-0.048*** (0.009)
Net Num. Targeted $\times$ Advice $\times$ Area Sales Volatility	-0.003 (0.005)	-0.032*** (0.005)	-0.032*** (0.005)
Net Num. Targeted $\times$ Advice $\times$ High Sales Volume Machine	-0.089*** (0.011)	-0.083*** (0.010)	-0.083*** (0.010)
Net Num. Targeted $\times$ Advice $\times$ High Sales Volatility Machine	0.060*** (0.010)	0.044*** (0.010)	0.045*** (0.010)
Net Num. Targeted $\times$ Integration	-0.045*** (0.008)	-0.069*** (0.008)	-0.068*** (0.008)
Net Num. Targeted $\times$ Integration $\times$ Regret	-0.031*** (0.008)	-0.036*** (0.007)	-0.036*** (0.007)
Net Num. Targeted $\times$ Integration $\times$ Delegation	0.010* (0.005)	0.009** (0.005)	0.010** (0.005)
Net Num. Targeted $\times$ Integration $\times$ Experience	-0.002 (0.009)	0.022*** (0.008)	0.021** (0.008)
Net Num. Targeted $\times$ Integration $\times$ Area Sales Volatility	0.025*** (0.005)	0.036*** (0.004)	0.036*** (0.004)
Net Num. Targeted $\times$ Integration $\times$ High Sales Volume Machine	0.099*** (0.011)	0.082*** (0.010)	0.083*** (0.010)
Net Num. Targeted $\times$ Integration $\times$ High Sales Volatility Machine	-0.040*** (0.010)	-0.019** (0.009)	-0.019** (0.009)
Num. Centralized	0.400*** (0.001)	0.243*** (0.002)	0.246*** (0.002)
R ²	0.478	0.565	0.566
Adj. R ²	0.478	0.564	0.565
Num. obs.	145688	145688	145688
RMSE	4.580	4.184	4.179

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are given in parentheses. Standard errors are clustered at the area level. Model 2 controls for product-specific fixed effects, and Model 3 controls for year-week-specific fixed effects.

Table A15: Do You Trust This Product Assortment Algorithm?

(%)	<i>All</i>		<i>Control</i>		<i>Treatment I</i>		<i>Treatment II</i>	
	<i>Before</i>	<i>After</i>	<i>Before</i>	<i>After</i>	<i>Before</i>	<i>After</i>	<i>Before</i>	<i>After</i>
Strongly agree	0.034	0.039	0.033	0.033	0.017	0.050	0.051	0.034
Agree	0.324	0.408	0.400	0.533	0.317	0.333	0.254	0.356
Disagree	0.363	0.330	0.333	0.317	0.383	0.367	0.373	0.305
Strongly Disagree	0.201	0.134	0.183	0.050	0.183	0.133	0.237	0.220
No Answer	0.078	0.089	0.050	0.067	0.100	0.117	0.085	0.085
N	179	179	60	60	60	60	59	59

¹ The unit of observations is an area.

Table A16: Do You Trust This Product Assortment Algorithm?: Ordered Probit

	<i>All</i>	<i>Control</i>	<i>Treatment I</i>	<i>Treatment II</i>	<i>All</i>
After	0.42** (0.21)	0.70* (0.36)	0.33 (0.36)	0.23 (0.35)	0.42** (0.21)
Experience ≥ 1 year					-0.94 (0.72)
AIC	773.23	251.01	255.78	271.71	773.45
BIC	788.40	261.92	266.48	282.44	792.41
Log Likelihood	-382.62	-121.50	-123.89	-131.85	-381.72
Deviance	765.23	243.01	247.78	263.71	763.45
Num. obs.	328	113	107	108	328

¹ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are given in parentheses.