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

Kohei Kawaguchi[*]

April 6, 2018

Abstract

This paper studies if and under what conditions workers follow an algorithm, in the context of experimentation with product assortments for beverage vending machines. I develop an algorithm to automate the human task, and use simulation to show that it increases revenue. I then conduct a field experiment where I introduce the algorithm. However, it fails to increase revenue, because workers are not willing to follow it. They follow only when their opinions are integrated. Workers having more regret from the last business semester follow the algorithmic advice more, and less when their opinions are integrated. [1][2]

[*]kkawaguchi@ust.hk, Department of Economics, School of Business and Management, Hong Kong University of Science and Technology. 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, and Andrew Ching 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.

[1] *JEL classification*: L20; L81; M12; M15; M20.

[2] *Keywords*: Automation; Experimentation; Learning; Product assortment problem; Multi-armed bandit problem; Algorithm aversion; Conflicts of interest; Productivity; Retail business; Beverage vending machines.

1 Introduction

Algorithms are taking over human decision-making in every business domain. However, full automation of a business activity is rather an exception. The exceptions include the Internet companies, in which a program also implements algorithmic decisions, and manufacturers, in which robots can work in a well-controlled environment. In several other situations, such as retail industries operating in a real-word shopping site, decisions regarding pricing, product assortment, and inventory control, must be physically implemented in an environment that is too complicated to automate. From these observations, we realize that we are not yet free from manpower.

With partial automation, a manager at the best can order workers to follow algorithmic advice. This paper studies if and under what conditions this strategy is effective. 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 stick to inferior human decisions, even if they understand that algorithmic decisions outperform them, is widely observed. This is known as *algorithm aversion* (Dietvorst, Simmons, & Massey, 2015). This type of aversion may prevent workers from following algorithms. Similarly, Hoffman, Kahn, & Li (2018) showed, in the context of hiring, that managers often overrule machine-based recommendation and result in worse outcomes. Dietvorst, Simmons, & Massey (2016) also proposed a way of reducing algorithm aversion. In a laboratory experiment, they demonstrated that people were willing to follow algorithms if only a small amount of control over the decision was allowed. It is important, however, to test these laboratory results with the field experiments in an actual business environment.

The cost of implementation is another source of conflicts of interest between the company and 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. For example, Atkin, Chaudhry, Chaudry, Khandelwal, & Verhoogen (2017) identified conflicts of interest as organizational barriers to technology adoption in the production of soccer ball in Pakistan.

My investigation takes the form of a field experiment in a large-scale beverage vending machine company operating in the Tokyo metropolitan area. Among the many tasks of this company, I focus on the task of experimentation with product assortments. This is because product assortment is one of the most important and computationally demanding decision variable in retail businesses as much as pricing. In the company, a worker, who are hired by a third-party subcontractor, is assigned to a narrowly defined area that consists of a few dozens of vending machines. A worker is responsible for machine maintenance and part of product assortment in the vending machines slots. The workers are incentivized to experiment with products and adapt the assortments to local demand. I can directly observe workers' responses when an algorithm takes over their decisions.

First, I develop an algorithm to automate the experimentation in the business setting of the company. It starts from a prior belief about the demand parameters of the products in an area during a business semester. The algorithm chooses the product assortments for each vending machine per area, on a weekly basis. At the end of the week, it observes the weekly data and updates the belief, continuing the process until the end of the semester. It balances the exploitation and exploration according to a certain rule. Because there is no readily available algorithm for this task, 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 my expriment. I demonstrate with a simulation 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 semester. In the experiment, I study 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 provide advice on the basis of the algorithm integrating the workers' forecasts for the demand of the new products into the prior belief of the algorithm. I tell the workers if their agorithm 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 shown to the workers. This indicates the existence of strong algorithm aversion or some other barriers to implementing the suggested product assortment decisions. The product assortments of the second treatment group correlated more with the algorithm's output than those of the first treatment group. Thus, integrating workers into the design of algorithms may resolve part of the problem. My findings are consistent with Dietvorst et al. (2015) and Dietvorst et al. (2016). The follow-up survey also confirms the importance of integrating workers' opinions, but the majority suggested that the increased burden of the operation prevented them from following the algorithm. This finding is consistent with Atkin et al. (2017).

Workers' reactions can be heterogeneous, depending on their ability, experience, and confidence, regarding their own product assortment decisions. In fact, a regression in which the treatment status interacts with the average regret of the workers in the same area during the previous semester shows that if the regret is higher i) workers follow the algorithmic advice more, and ii) they follow less when their beliefs are integrated into the algorithm. This result has an important management implication. Algorithmic advice selectively affects worker behavior, and self-selection works in a way in which worse workers conform more.

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

For this research, I build an algorithm to automate the experimentation of product assortment decisions in the beverage vending machine business based, on the basis of the ideas developped 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 & Li (2011), and May, Korda, Lee, & Leslie (2012), demonstrated the empirical efficacy of TS vis simulations. Agrawal & Goyal (2012) showed that the TS algorithm achieved logarithmic expected regret for the stochastic multi-armed bandit problem. Caro & Gallien (2007) and Rusmevichientong, Max Shen, & Shmoys (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 & Gallien (2007)'s policy exploited a specific form of demand that had no substitution across products. Thus, it was not empirically appropriate.

In laboratory experiments, Dietvorst et al. (2015) found that humans tend to possess an algorithm aversion, and Dietvorst et al. (2016) found that the algorithm aversion could be eased by ceding a small amount of control. This paper supports these findings with a field experiment. There is also a literature that theorizes the value of authority and control (Bartling, Fehr, & Herz, 2014; Fehr, Herz, & Wilkening, 2013; Owens, Grossman, & Fackler, 2014). My paper provides a motivating example, b it dutoes not directly test the mechanisms proposed in those papers.

Several studies such as those by Israel (2005), Erdem, Keane, Öncü, & Strebel (2005), Narayanan, Manchanda, & Chintagunta (2005), Chernew, Gowrisankaran, & Scanlon (2008), Erdem, Keane, & Sun (2008), Narayanan & Manchanda (2009), A. T. Ching (2010), Anand & Shachar (2011), Szymanowski & Gijsbrechts (2012), and A. T. Ching, Erdem, & Keane (2013), used consumer choice data to estimate consumer learning models and examine its implications for marketing and policy interventions. A similar approch was employed to study the efficacy of physician's learning by Coscelli & Shum (2004), Crawford & Shum (2005), and A. Ching & Ishihara (2010). There are several papers that theoretically studied the experimentation by firms, including Rob (1991), Mirman, Samuelson, & Urbano (1993), Creane (1994), Raman & Chatterjee (1995), and Hitsch (2006) in a monopoly market, and Aghion, Espinosa, & Jullien (1993) and Harrington (1995) in a duopoly market, and Harpaz, Lee, & Winkler (1982), Bertocchi & Spagat (1998), and Bergemann & Valimaki (2000) in a more competitive market. They mostly studied price experimentation, whereas the current paper is about experimentation with product assortment deisions. Additionally, my paper is empirical. Bolton & Harris (1999), Keller, Rady, Cripps, & Cripps1 (2005), and Keller & 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,[3] 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.

[3] http://www.jvma.or.jp/index.html.

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.[4]

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 business, the company operates a retail business inside stations to exploit the substantial number of passengers who regularly use the transportation service.[5] 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

[4]The 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 (R. Deneckere, Marvel, & Peck, 1996; Raymond Deneckere, Marvel, & Peck, 1997).

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

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 semester. 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 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, Takano, Nakajima, Go, & Yasuoka, 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.[6] Because workers regularly visit vending machines regardless of product changes, marginal costs of changing products atop driving, visiting and refilling are not high.

2.3 Product Assortment Policy

The business plan is determined for each semester. 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 semester. Until then, the product cannot 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

[6]Principal 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.

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 semesters. It also change 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 semester. 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.

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 semester is indexed by $i \in \mathcal{I}$. In each local area and semester, 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 $\mathcal{J}_g$. 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 chocie specification is made for a practical reason. In the following analysis, I solve the optimal product assortments given a demand function for many occurr. When the demand function is a nested-logit form, I can utilize the algorithm of Gallego & Topaloglu (2014) to solve a static product assortment problem with a group-by-group shelf restriction. One may think that the demand function should be a mixed-logit form, but Rusmevichientong, Shmoys, Tong, & Topaloglu (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 $\gamma_i \to 1$, 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. We 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:

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

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

$$v_{ijt} = \xi_{ij} + \beta_{ig} \times Temperature_{it},$$

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 semester.

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

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

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 semester 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 semester and compute the *ex post* belief mean.[7]

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.

[7]If a product is not observed in a local area in a semester, 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 semester.

3.2 Estimation Results

The model fits the data well. Figure 4 displays the actual versus the predicted 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 plot. The R^2 of the model is 0.586. The summary statistics of the *ex post* belief across areas and semesters are shown in Tables 4, 5, and 6. 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 semesters. 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 semesters. Table 6 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 semesters 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 $\boldsymbol{d}_{ikt}$ be the vector of product assortments at vending machine k in area-semester i in week t. Then, the revenue of product j at the vending machine under the product assortments is:

$$
\begin{aligned}
r_{ijkt} &= p_j \frac{\exp(v_{ijt}/\gamma_i) d_{ijkt}}{\sum_{l \in J_g} \exp(v_{ilt}/\gamma_i) d_{ilkt}} \frac{\left(\sum_{l \in J_g} \exp(v_{ilt}/\gamma_i) d_{ilkt} \right)^{\gamma_i}}{1 + \sum_{m \in \mathcal{G}} \left(\sum_{l \in J_m} \exp(v_{ilt}/\gamma_i) d_{ilkt} \right)^{\gamma_i}} \exp(\sigma_i \epsilon_{ijkt}) \\
&\equiv r_{ijt}(\boldsymbol{d}_{ikt}, \boldsymbol{\xi}_i, \boldsymbol{\beta}_i, \gamma_i) \exp(\sigma_i \epsilon_{ijkt}),
\end{aligned}
$$

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

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

I design an agent that sequentially decides product assortments $\boldsymbol{d}_{it}$ while learning the demand from the past sales data $\boldsymbol{r}_{i,1:t-1} = (r_{i1}, \cdots, r_{i,t-1})'$. Let $F_{i,t+1}$ be the belief of the agent over the local demand parameters at the end of period t, and let $B(F_{i,t+1}; F_{it}, \boldsymbol{d}_{it})$ 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 $\boldsymbol{\epsilon}_{it}$.

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

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

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:

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

subject to:

$$
\sum_{j \in \mathcal{J}_g} d_{ijkt}(F_{it}) = c_{igkt}, k \in \mathcal{K}, g \in \mathcal{G},
$$

$$
d_{ijkt} \geq \underline{d}_{ijkt}(F_{it}), j \in \mathcal{J}, k \in \mathcal{K}.
$$

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

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 & 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. To the best of my knowledge, there is no other study on this kind of setting. Caro & 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 & 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 & Li (2011), May et al. (2012) have demonstrated the empirical efficacy of TS using simulations. Agrawal & Goyal (2012) showed that the TS achieves logarithmic expected regret for the stochastic multi-armed bandit problem. Ortega & Braun (2010), who called the algorithm the "Bayesian control rule"", provides a high-level condition for the convergence of the policy toward the ex-post optimal policy.

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 $\boldsymbol{\xi}^*_{ik}, \boldsymbol{\beta}^*_{ik}, \gamma^*_{ik}$. ii) Then, at vending machine k, the agent chooses a product assortment that solves:

$$\max_{\boldsymbol{d}_{ikt}} \sum_{j \in \mathcal{J}} r_{ijt}(\boldsymbol{d}_{ikt}, \boldsymbol{\xi}^*_{ik}, \boldsymbol{\beta}^*_{ik}, \gamma^*_{ik}),$$

from a feasible set of assortment. This is the optimal product assortment when the true demand parameter is $\boldsymbol{\xi}^*_{ik}, \boldsymbol{\beta}^*_{ik}, \gamma^*_{ik}$, 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 semester. 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 & Topaloglu (2014) to solve a static product assortment problem with a group-by-group shelf 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:

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

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 $B(F_{i,t+1}; F_{it}, \boldsymbol{d}_{it})$, 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 semester, only the product-specific components of the demand function, $\boldsymbol{\xi}_i$, are learned online during the semester. The group-specific components $\boldsymbol{\beta}_i$, γ_i, and σ_i are not learned. Instead, I use the average of the estimates across semesters 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 semester is set at the *ex post* belief for them in the last semester 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, $B(F_{i,t+1}; F_{it}, \boldsymbol{d}_{it})$, for $\boldsymbol{\xi}_i$ 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 othe 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 are defined by:

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

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

Table 7 and Figure 5 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 actuals. 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 $\tau = Actual, TS$,

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

with and without the inequality restrictions H_0: $\gamma_l \geq 0$ 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. Hence, the null hypothesis that the success rates of TS product assortments is rejected.

I also compare the regret under TS to actual product assortments. Table 8 and Figure 6 show the mean and 99% confidence intervals of the regret under TS and actual product assortments over time. 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. 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 experimentat 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. We 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 the existence of a strong algorithm aversion. If the correlation is higher in treatment II than in treatment I, the integration of the workers' forecasts into the algorithm reduces the algorithm aversion.

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. 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. I did not provide information about the other areas. However, I could not prohibit the workers from sharing the information. 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, and I do not appear to the workers. 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

I conducted a survey asking workers their forecast for the products planned to be introduced during the coming semester, Fall/Winter 2016, before the products were introduced. During this semester, 24 distinct products were introduced. Some were introduced in early fall and some in late fall. The survey was conducted in August, after the company prepared its plan for the coming semester.

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, if a worker predicts that the product in the green tea category sells 120 units in the week starting October 17 where the average temperature is 16°C, the average sales units of green tea is 100. Thus, the worker is advised to write their prediction as 1.2. 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 semester. 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 9, 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 semester 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 semester 2016, denoted by $\overline{\xi}_{ig}$. 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 semester 2016.

To see how these values, $\xi_{ij}^{implied}$ and $\overline{\xi}_{ig}$, are informative, I regressed the *ex post* belief of ξ_{ij} from the end of the Fall/Winter semester 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 10 shows the regression results, when only $\overline{\xi}_{ig}$ is included, only $\xi_{ij}^{implied}$ is included, and when both $\xi_{ij}^{implied}$ and $\overline{\xi}_{ig}$ are included. All coefficients except for the constant, are positive and statistically significant. However, the size of the coefficient is larger for $\overline{\xi}_{ig}$ 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 $\overline{\xi}_{ig}$ 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 Experiment Results

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

$$
\begin{aligned}
ActualNumSlot_{ijt} = {} & \beta_1 NumNetTargetedSlot_{ijt} \\
& + \beta_2 NumNetTargetedSlot_{ijt} \times Advice_i \\
& + \beta_3 NumNetTargetedSlot_{ijt} \times Integration_i \\
& + \beta_4 NumCentralizedSlot_{ijt} \\
& + Controls + \epsilon_{ijt},
\end{aligned}
$$

where i is an area, j is a product, t is a target week, $ActualNumSlot_{ijt}$ is the number of actual slots of 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 prior belief of the workers. Our main parameters of interest are β_2 and β_3, which capture the difference in correlation with the algorithmic advice between treatment statuses. I estimated model controlling for several sets of unobserved fixed effects for a robustness check, but controlling for them was not necessary for the validity of the current analysis.

Table 11 shows the estimation results. (i) $Advice_i$ did not significantly affect the product assortment decision of workers, and (ii) the integration of workers' belief statistically significantly increased the conformity of workers with algorithmic advice. This implies that (i) there was a strong algorithm aversion and that (ii) the integration of the workers into the algorithm design process reduced the extent of the algorithm aversion. 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 with the control group. The correlation, however, was 0.118 in the control group. The change in the extent of correlation can be regarded large enough relative to the baseline correlation (13.011%) and as the effect of just integrating the workers' prior belief into the algorithm. However, at the absolute level, the effect was too small to cause a meaningful change in the product assortment decisions of the workers. In fact, another regression showed that there were no statistically significant differences in the sales level and the year-on-year sales growth rate between the control and the treatment groups.

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 provided instructions in treatment I were more similar to the idea of the workers, 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 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 prior belief quickly diminish once actual sales are observed. Therefore, the first channel will likely dominate.

The reaction of workers is expected to be heterogeneous, depending on the ability, experience, and confidence with own product assortment decisions. To see this, I computed the difference in the weekly regret of actual and TS product assortment decisions of each area during the Spring/Summer 2016. I took the average of the weekly regret for each area. I used this measure, $Regret_i$, as the potential benefit of using algorithmic advice for workers in each area. Although the identity of workers in an area can change across semesters, the change is not substantial within the same year, according to the managers of the company. Table 12 shows the estimation results when I interact $Regret_i$ 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) follow less when their belief is integrated into the algorithm. This result has an important management implication: algorithmic advice selectively affects the behavior of workers, and the self-selection works in a way in which worse workers conform more.

5.4 Follow-up Survey

At the end of March 2017 after the experiment, I conducted a review survey. Because this survey was conducted under the name of the company, the response rate was 100%.

In the survey, I 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 either "strongly agree", "agree", "disagree" or "strongly disagree". Table 13 summarizes their answers. Most workers disagreed both before and after the experiment. However, overall, the proportion of workers who agreed increased, and those who disagreed decreased after the experiment. Before the experiment, 0.363 disagreed and 0.201 strongly disagreed. However, after the experiment 0.330 disagreed and 0.134 strongly disagreed. Before the experiment, 0.324 agreed and 0.034 strongly agreed. However, after the experiment 0.408 agreed and 0.039 strongly agreed. This pattern is universal across 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, indicating that the average experience of workers in the area exceeds 1 year. The regression results are shown in Table 14. 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, trust only increases in the control group in which the 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 is because people are less tolerant of machines' errors than human errors. The fifth column uses the entire sample and includes 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.

Second, using multple choice, I asked "What kind of improvement is needed for you to follow the algorithm?". The answers were reliable demand prediction, integration of workers' prediction, integration of managers' prediction, detailed explanation of the algorithm, coordination with manager's order, reduced burden of the work, machine-level instruction, inside-station-level instruction, station-level instruction, mobile interface, visual interface, coordination with product abolition schedule, and compensation for loss. The workers could list as many items as they wanted in their answers. I summarized the reply in unweighted and weighted manners. In the unweighted result, for each item, I counted the number of areas that include the item in the answer. In the weighted result, for each item, if the answer of an area included the item, I divided 1 by the number of items included in the answer of the area. Then, I summed up the weighted numbers across areas. Table 15 summarizes the result. The patterns are qualitatively the same between weighted and unweighted summaries, and across treatment statuses.

The required improvement is the reduced burden of the work. Overall, scores are 101 (unweighted) and 32.2 (weighted). This highlights the friction that may only exist in the *offline* business. In the case of the *online* business, decisions made by an algorithm can be automatically implemented, once the code is written. However, in the case of the *offline* business, such as a retail business with physical stores, decisions made by an algorithm have to be physically implemented. In the future, or in some field, this may be done by robots, but currently in many fields, this must be done by humans. Unless these physical implementation costs are removed, the introduction of algorithms into businesses cannot achieve the aimed efficiency. That substantial share of workers requiring coordination with a product abolition schedule is related to the burden. Because beverages are seasonal products, some products are planned to be abolished at some point in time. This requires workers to resolve 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 of the baseline demand prediction. The scores are 76 (unweighted) and 21.1 (weighted). Relatedly, a substantial proportion of workers required more detailed explanation of the algorithm. In addition to improving the algorithm, establishing better communications 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 the instruction 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 simplify the content by themselves.

The extent of the integration of workers' predictions was also substantial. The scores are 52 (unweighted) and 15.4 (weighted). Interestingly, the workers did not demand to integrate 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 we can easily improve upon introducing an algorithm into business. The results from the last section support the notion that this facilitates the use of algorithms by workers.

In Table 16, I report the estimation results of the effect of algorithmic advice, where some of the answers to the follow-up survey interact. In the regression, $Trust_i$ assigns values from 4 to 1 to answers "strongly agree" to "strongly disagree", respectively, to the question "Do you trust this product assortment algorithm?" $Understanding_i$ assigns values from 4 to 1 to answers "strongly agree" to "strongly disagree" to question "Do you understand this product assortment algorithm?" $Experience_i$ assigns values from 4 to 1 to answers "More than 10 years”, “More than 5 years no greater than 10 years”, “More than 1 years no greater than 5 years”, and “No greater than 1 years”, respectively, to the question “How many years did you work for this business?”. The result shows that these answers have little explanatory power to the heterogeneous reaction of workers to algorithmic advice.

6 Conclusions

In this paper, I studied if and under what conditions workers would follow an algorithm, in the context of experimentation in product assortment for a beverage vending machine business. First, I developed an algorithm to automate the task of experimenting products in the current business setting, demonstrating that the algorithm can increase revenue, at least in a simulation. The simulation exercise showed that the potential impact of automation was significant in the current business setting. However, in the field experiment, I found that the algorithm failed to increase the revenue. The reason was that workers did not follow the algorithm. The follow-up survey hinted that conflict of interests, owing to the cost of implementing the suggested product assortments, the reliability of the baseline demand estimation, and the feeling that their opinions should be integrated more were major factors that discouraged workers from following the algorithm. The field experiment also revealed that the integration of workers’ forecasts into the algorithm could slightly encourage workers to follow. The results also showed that workers in an area having more regret in the previous semestors tend to follow the algorithmic advice more, and less when their opinions are integrated.

Acemoglu & 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 was a productivity effect. In the current setting, if the company can 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. My paper highlighted these issues in the context of experimentation in product assortment decisions.

A limitation of the research is that this was 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. My analysis did not consider joint pricing and product assortment decisions because pricing was not an issue for the current setting. It will be possible to extend my framework to such 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.

Reference

Acemoglu, D., & Restrepo, P. (2017). The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment. *SSRN Electronic Journal*.

Aghion, P., Espinosa, M. P., & Jullien, B. (1993). Dynamic Duopoly with Learning through Market Experimentation. *Economic Theory*, *3*(3), 517–539.

Agrawal, S., & Goyal, N. (2012). 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, pp. 1–39).

Anand, B. N., & Shachar, R. (2011). Advertising, the Matchmaker. *The RAND Journal of Economics*, *42*(2), 205–245.

Atkin, D., Chaudhry, A., Chaudry, S., Khandelwal, A. K., & Verhoogen, E. (2017). Organizational Barriers to Technology Adoption: Evidence from Soccer-ball Producers in Pakistan. *Quarterly Journal of Economics*, *132*(3), 1101–1164.

Bartling, B., Fehr, E., & Herz, H. (2014). The Intrinsic Value of Decision Rights. *Econometrica*, *82*(6), 2005–2039.

Bergemann, D., & Valimaki, J. (2000). Experimentation in Markets. *Review of Economic Studies*, *67*(2), 213–234.

Bertocchi, G., & Spagat, M. (1998). Growth under Uncertainty with Experimentation. *Journal of Economic Dynamics and Control*, *23*(2), 209–231.

Bolton, P., & Harris, C. (1999). Strategic Experimentation. *Source: Econometrica Econometrica*, *67*(2), 349–374.

Caro, F., & Gallien, J. (2007). Dynamic Assortment with Demand Learning for Seasonal Consumer Goods. *Management Science*, *53*(2), 276–292.

Chapelle, O., & Li, L. (2011). An Empirical Evaluation of Thompson Sampling. In *Advances in neural information processing systems (nips)*.

Chernew, M., Gowrisankaran, G., & Scanlon, D. P. (2008). Learning and the Value of Information: Evidence from Health Plan Report Cards. *Journal of Econometrics*, *144*(1), 156–174.

Ching, A. T. (2010). Consumer Learning and Heterogeneity: Dynamics of Demand for Prescription drugs after Patent Expiration. *International Journal of Industrial Organization*, *28*, 619–638.

Ching, A. T., Erdem, T., & Keane, M. P. (2013). Learning Models: An Assessment of Progress, Challenges, and New Developments. *Marketing Science*, *32*(6).

Ching, A., & Ishihara, M. (2010). The Effects of Detailing on Prescribing Decisions under Quality Uncertainty. *Quantitative Marketing and Economics*, *8*(2), 123–165.

Coscelli, A., & Shum, M. (2004). An Empirical Model of Learning and Patient Spillovers in New Drug Entry. *Journal of Econometrics*, *122*(2), 213–246.

Crawford, G. S., & Shum, M. (2005). Uncertainty and Learning in Pharmaceutical Demand. *Econometrica*, *73*(4), 1137–1173.

Creane, A. (1994). Experimentation with Heteroskedastic Noise. *Economic Theory*, *4*(2), 275–286.

Deneckere, R., Marvel, H. P., & Peck, J. (1996). Demand Uncertainty, Inventories, and Resale Price Maintenance. *The Quarterly Journal of Economics*, *111*(3), 885–913.

Deneckere, R., Marvel, H. P., & Peck, J. (1997). Demand Uncertainty and Price Maintenance: Markdowns as Destructive Competition. *Source: The American Economic Review*, *87*(4), 619–641.

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err. *Journal of Experimental Psychology: General*, *144*(1), 114–126.

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2016). Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them. *Management Science*.

Erdem, T., Keane, M. P., & Sun, B. (2008). A Dynamic Model of Brand Choice When

Price and Advertising Signal Product Quality. *Marketing Science*, *27*(6).

Erdem, T., Keane, M. P., Öncü, T. S., & Strebel, J. (2005). Learning About Computers: An Analysis of Information Search and Technology Choice. *Quantitative Marketing and Economics*, *3*(3), 207–247.

Fehr, E., Herz, H., & Wilkening, T. (2013). The Lure of Authority: Motivation and Incentive Effects of Power. *American Economic Review*, *103*(4), 1325–1359.

Gallego, G., & Topaloglu, H. (2014). Constrained Assortment Optimization for the Nested Logit Model. *Management Science*, *60*(10).

Gittins, J., & Jones, D. (1974). A Dynamic Allocation Index for the Sequential Design of Experiments. In J. Gani (Ed.), *Progress in statistics* (pp. 241–266). Amsterdam: North-Holland.

Harpaz, G., Lee, W. Y., & Winkler, R. L. (1982). Learning, Experimentation, and the Optimal Output Decisions of a Competitive Firm. *Management Science*, *28*(6), 589–603.

Harrington, J. E., Jr. (1995). Experimentation and Learning in a Differentiated-Products Duopoly. *Journal of Economic Theory*, *66*(1), 275–288.

Hitsch, G. J. (2006). An Empirical Model of Optimal Dynamic Product Launch and Exit Under Demand Uncertainty. *Marketing Science*, *25*(1), 25–50.

Hoffman, M., Kahn, L. B., & Li, D. (2018). Discretion in Hiring. *The Quarterly Journal of Economics*, *133*(2), 765–800.

Israel, M. (2005). Services as Experience Goods: An Empirical Examination of Consumer Learning in Automobile Insurance. *American Economic Review*, *95*(5), 1444–1463.

Keller, G., & Rady, S. (2010). Strategic Experimentation with Poisson Bandits. *Theoretical Economics*, *5*(2), 275–311.

Keller, G., Rady, S., Cripps, M., & Cripps1, M. (2005). Strategic Experimentation with Exponential Bandits. *Econometrica*, *73*(1), 39–68.

May, B. C., Korda, N., Lee, A., & Leslie, D. S. (2012). Optimistic Bayesian Sampling in Contextual-Bandit Problems. *Journal of Machine Learning Research*, 2069–2106.

Mirman, L. J., Samuelson, L., & Urbano, A. (1993). Monopoly Experimentation. *International Economic Review*, *34*(3), 549.

Morita, T., Takano, Y., Nakajima, T., Go, H., & Yasuoka, H. (2012). *New Marketing Strategy using the Electronic Money at Beverage Vending Machines (in Japanese).*

Narayanan, S., & Manchanda, P. (2009). Heterogeneous Learning and the Targeting of Marketing Communication for New Products. *Marketing Science*, *28*(3).

Narayanan, S., Manchanda, P., & Chintagunta, P. K. (2005). Temporal Differences in the Role of Marketing Communication in New Product Categories. *Journal of Marketing Research*, *42*(3), 278–290.

Ortega, P. A., & Braun, D. A. (2010). Convergence of the Bayesian Control Rule.

Owens, D., Grossman, Z., & Fackler, R. (2014). The Control Premium: A Preference for Payoff Autonomy. *American Economic Journal*, *6*(4), 138–161.

Raman, K., & Chatterjee, R. (1995). Optimal Monopolist Pricing Under Demand Uncertainty in Dynamic Markets. *Management Science*, *41*(1), 144–162.

Rob, R. (1991). Learning and Capacity Expansion under Demand Uncertainty. *The Review of Economic Studies*, *58*(4), 655.

Rusmevichientong, P., Max Shen, Z.-J., & Shmoys, D. B. (2010). Dynamic Assortment Optimization with a Multinomial Logit Choice Model and Capacity Constraint. *Operations Research*, *58*(6), 1666–1680.

Rusmevichientong, P., Shmoys, D., Tong, C., & Topaloglu, H. (2014). Assortment Optimization under the Multinomial Logit Model with Random Choice Parameters. *Production and Operations Management*, *23*(11), 2023–2039.

Scott, S. L. (2010). A Modern Bayesian Look at the Multi-armed Bandit. *Applied Stochastic Models in Business and Industry*, *26*, 639–658.

Szymanowski, M., & Gijsbrechts, E. (2012). Consumption-Based Cross-Brand Learning:

Are Private Labels Really Private? *Journal of Marketing Research*, *49*(2), 231–246.

Thompson, W. (1933). On the Likelihood that One Unknown Probability Exceeds in View of the Evidence of Two Samples. *Biometrika*, *25*(3/4), 285–294.

Wang, R. (2012, November). Capacitated Assortment and Price Optimization under the Multinomial Logit Model. North-Holland.

List of Figures

Figure 1: Railway Network of JRE

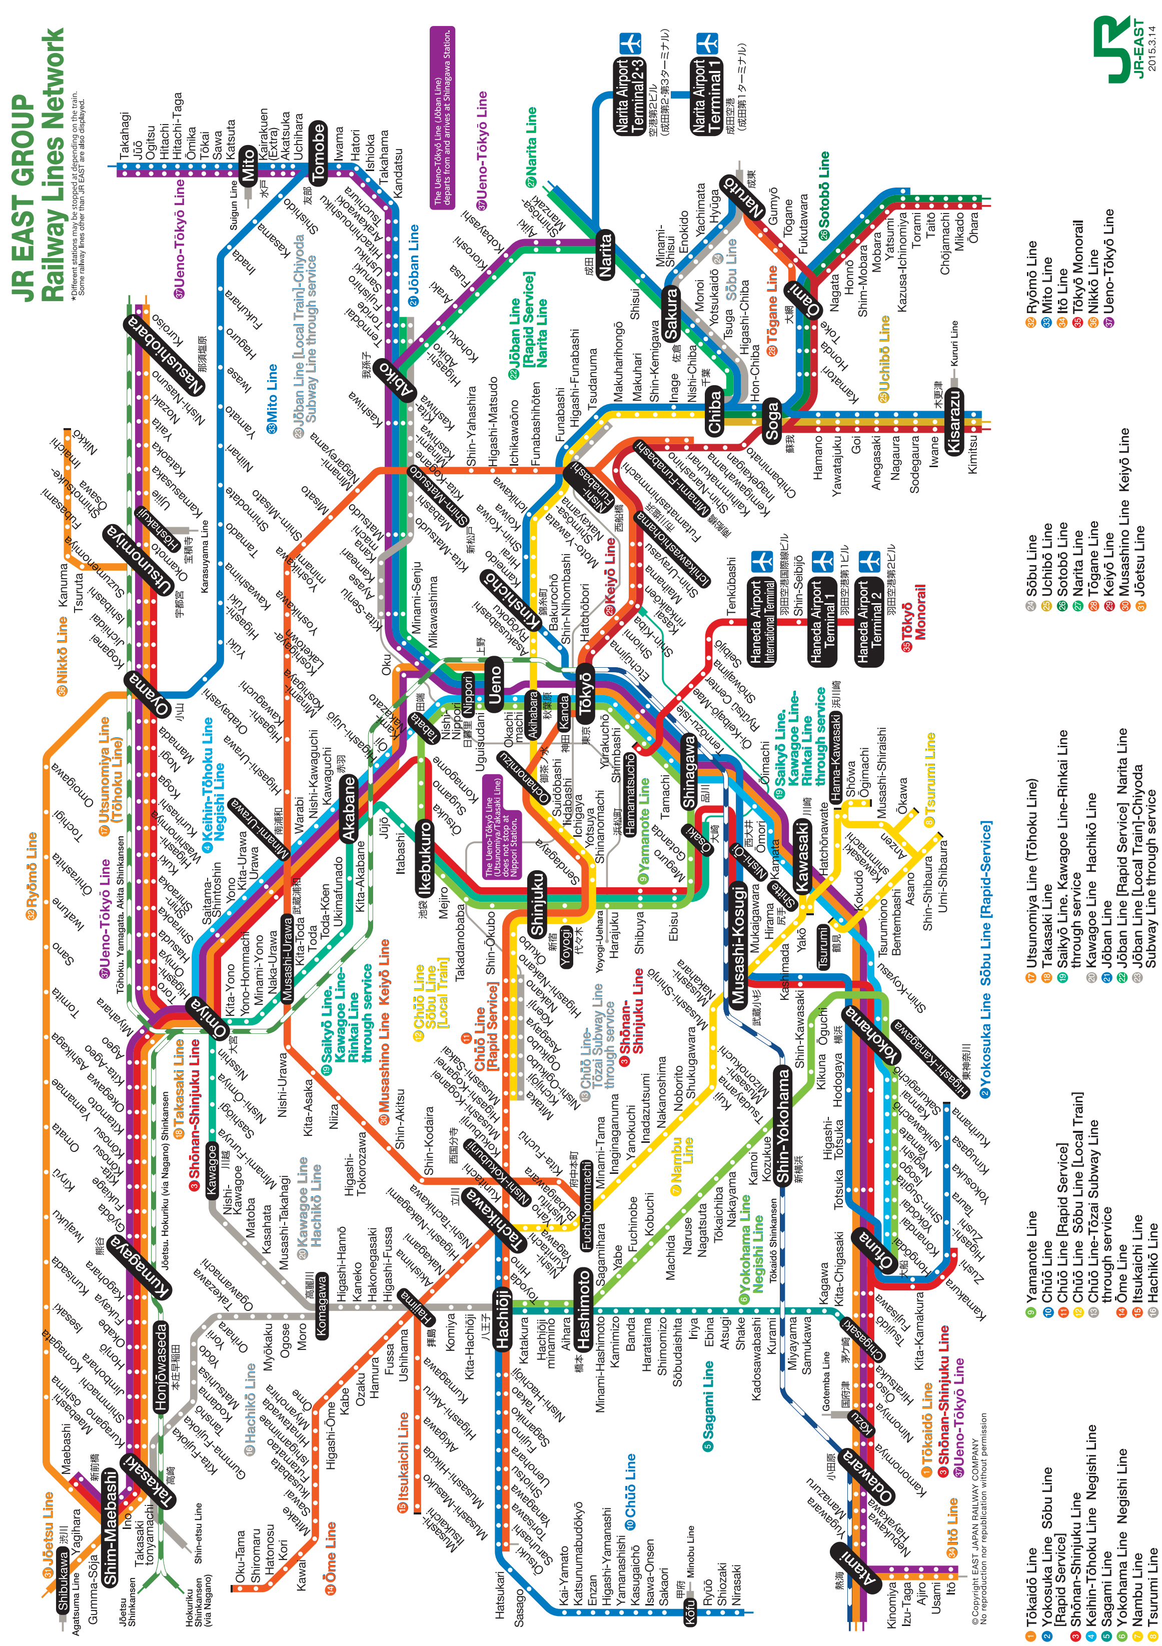

Figure 2: A Beverage Vending Machine of JRE

Figure 3: A Vending Machine Operation Worker of JRE

Figure 4: The Fit of the Demand Function

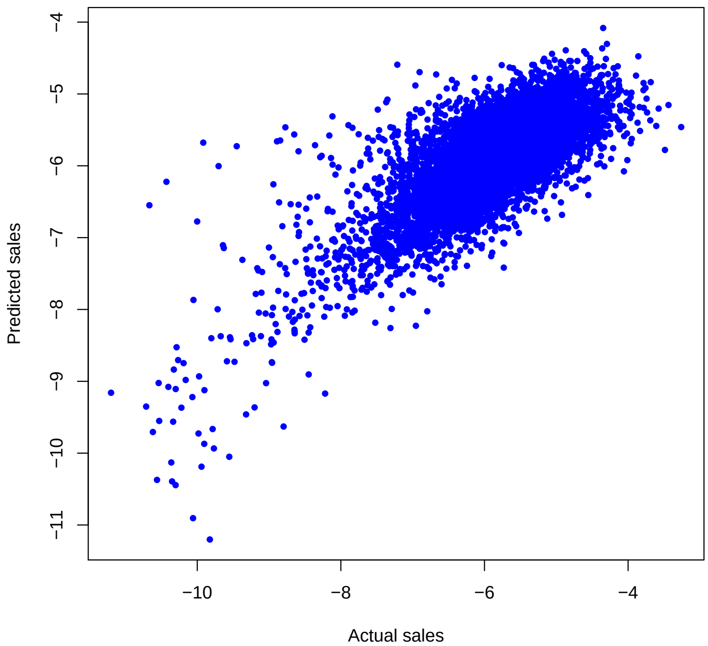

[1] The scale is the log difference from the outside share. Because the sample size is huge, I randomly sampled 10,000 observations to display the plot.

Figure 5: Mean TS and Actual Success Rates Over Time

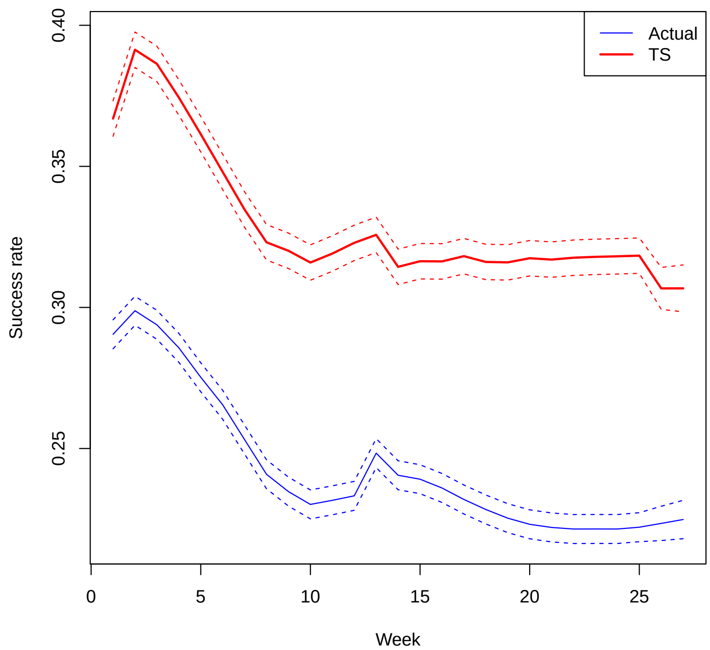

[1] The dashed lines are 99% confidence intervals of the mean success rates.

Figure 6: Mean TS and Actual Regret Over Time

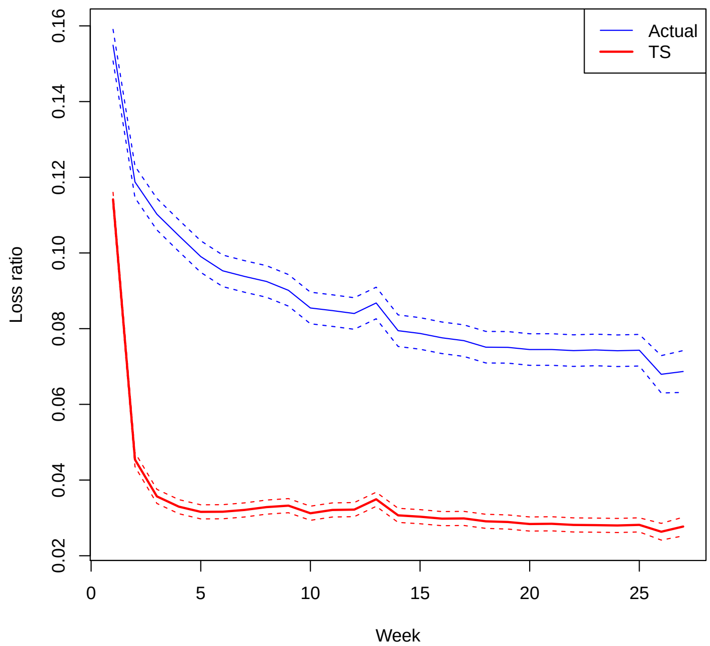

[1] The dashed lines are 99% confidence intervals of the mean regrets.

List of Tables

Table 1: Number of Unique Values

Semester	*Product*	*Machine*	*Station*	*Area*	*Subcontractor*
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

	N	*Mean*	*Sd*	*Min*	*Max*	*Q1*	*Q2*	*Q3*
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

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 between 2013 Summer/Spring and 2016 Summer/Spring. It only counts the products on delegated slots.

Table 4: *Ex Post* Belief About γ and σ^2

Parameter	*N*	*Mean*	*Sd*	*Min*	*Max*	*Q1*	*Q2*	*Q3*
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 semesters, and N is the number of relevant areas and semesters.

Table 5: *Ex Post* Belief About β

Group	*N*	*Mean*	*Sd*	*Min*	*Max*	*Q1*	*Q2*	*Q3*
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
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 semesters, and N is the number of relevant areas and semesters.

Table 6: *Ex Post* Belief About ξ

Product	*N*	*Mean*	*Sd*	*Min*	*Max*	*Q1*	*Q2*	*Q3*
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 semesters, and N is the number of relevant areas and semesters. The percentile is among products that are estimated more than 100 areas and semesters. I cannot provide the product names because of confidentiality reasons.

Table 7: Mean TS and Actual Success Rates Over Time

	TS	*Actual*
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 8: Mean Regret under TS and Actual Product Assortments Over Time

	TS	*Actual*
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 9: Prior Belief for New Products

Setting	*N*	*Mean*	*Sd*	*Min*	*Max*	*Q1*	*Q2*	*Q3*
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 10: Predictive Power of Worker's Prior Belief for New Goods

	(1)	(2)	(3)
Constant	-2.61^{***}	-5.87^{***}	-2.56^{***}
	(0.02)	(0.01)	(0.02)
Average Belief	0.88^{***}		0.86^{***}
	(0.00)		(0.00)
Worker Belief		0.29^{***}	0.03^{***}
		(0.00)	(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 in parentheses.

Table 11: The Effect of Algorithmic Advice on the Assortment Decisions

	(1)	(2)	(3)
Net Num. Targeted	0.33^{***}	0.11^{***}	0.12^{***}
	(0.01)	(0.01)	(0.01)
Net Num. Targeted × Advice	-0.01	0.00	0.00
	(0.01)	(0.01)	(0.01)
Net Num. Targeted × Integration	0.02^{**}	0.01^{**}	0.01^{**}
	(0.01)	(0.01)	(0.01)
Num. Centralized	1.09^{***}	0.84^{***}	0.84^{***}
	(0.00)	(0.00)	(0.00)
R^2	0.74	0.84	0.84
Adj. R^2	0.74	0.84	0.84
Num. obs.	74029	74029	74029
RMSE	5.84	4.61	4.59

[1] $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.10$. Standard errors are in parentheses. Model 2 in addition controls for product-specific fixed effects. Model 3 in addition controls for year-week-specific fixed effects.

Table 12: The Effect of Algorithmic Advice on the Assortment Decisions

	(1)	(2)	(3)
Net Num. Targeted	0.34^{***}	0.12^{***}	0.12^{***}
	(0.01)	(0.01)	(0.01)
Net Num. Targeted × Regret	-0.35	-0.14	-0.16
	(0.34)	(0.27)	(0.27)
Net Num. Targeted × Advice	-0.06^{***}	-0.04^{***}	-0.04^{***}
	(0.02)	(0.01)	(0.01)
Net Num. Targeted × Advice × Regret	1.51^{***}	1.26^{***}	1.26^{***}
	(0.44)	(0.35)	(0.35)
Net Num. Targeted × Integration	0.08^{***}	0.08^{***}	0.08^{***}
	(0.02)	(0.01)	(0.01)
Net Num. Targeted × Integration × Regret	-1.75^{***}	-1.77^{***}	-1.77^{***}
	(0.37)	(0.29)	(0.29)
Num. Centralized	1.09^{***}	0.84^{***}	0.84^{***}
	(0.00)	(0.00)	(0.00)
R^2	0.74	0.84	0.84
Adj. R^2	0.74	0.84	0.84
Num. obs.	74029	74029	74029
RMSE	5.84	4.61	4.59

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

Table 13: Do You Trust This Product Assortment Algorithm?

(%)	*All*		*Control*		*Treatment I*		*Treatment II*	
	Before	*After*	*Before*	*After*	*Before*	*After*	*Before*	*After*
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 14: Do You Trust This Product Assortment Algorithm?: Ordered Probit

	All	*Control*	*Treatment I*	*Treatment II*	*All*
After	0.42^{**}	0.70^{*}	0.33	0.23	0.42^{**}
	(0.21)	(0.36)	(0.36)	(0.35)	(0.21)
Experience $\geq$ 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 in parentheses.

Table 15: What are the Barriers? (Multiple Answers)

(a) Unweighted

Answer	*All*	*Control*	*Treatment I*	*Treatment II*
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

Answer	*All*	*Control*	*Treatment I*	*Treatment II*
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.

Table 16: The Effect of Algorithmic Advice on the Assortment Decisions

	(1)	(2)	(3)
Net Num. Targeted	0.29^{***}	0.14^{***}	0.14^{***}
	(0.04)	(0.03)	(0.03)
Net Num. Targeted × Regret	-0.63	-0.73^{**}	-0.73^{**}
	(0.39)	(0.31)	(0.31)
Net Num. Targeted × Trust	0.00	0.01^{**}	0.01^{**}
	(0.01)	(0.01)	(0.01)
Net Num. Targeted × Understanding	0.00	-0.00	-0.00
	(0.01)	(0.01)	(0.01)
Net Num. Targeted × Experience	0.01	-0.01	-0.01
	(0.01)	(0.01)	(0.01)
Net Num. Targeted × Advice	-0.10^{*}	-0.04	-0.04
	(0.05)	(0.04)	(0.04)
Net Num. Targeted × Advice × Regret	1.98^{***}	1.25^{***}	1.26^{***}
	(0.51)	(0.41)	(0.40)
Net Num. Targeted × Advice × Trust	0.00	-0.01	-0.01
	(0.01)	(0.01)	(0.01)
Net Num. Targeted × Advice × Experience	0.00	0.01	0.01
	(0.01)	(0.01)	(0.01)
Net Num. Targeted × Integration	0.05	-0.01	-0.00
	(0.05)	(0.04)	(0.04)
Net Num. Targeted × Integration × Regret	-1.95^{***}	-0.95^{***}	-0.97^{***}
	(0.42)	(0.33)	(0.33)
Net Num. Targeted × Integration × Trust	0.00	0.00	0.00
	(0.01)	(0.01)	(0.01)
Net Num. Targeted × Integration × Understanding	0.02	0.01	0.01
	(0.01)	(0.01)	(0.01)
Net Num. Targeted × Integration × Experience	-0.01	0.01	0.01
	(0.01)	(0.01)	(0.01)
Num. Centralized	1.09^{***}	0.83^{***}	0.84^{***}
	(0.00)	(0.00)	(0.00)
R^2	0.74	0.84	0.84
Adj. R^2	0.74	0.84	0.84
Num. obs.	65873	65873	65873
RMSE	5.85	4.60	4.58

[1] $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.10$. Standard errors are in parentheses. Standard errors are clustered at the area level. Model 2 in addition controls for product-specific fixed effects. Model 3 in addition controls for year-week-specific fixed effects. $Regret_i$ is the average of actual regret minus TS regret. $Trust_i$ assigns values from 4 to 1 to answers "strongly agree" to "strongly disagree" to the question "Do you trust this product assortment algorithm?" $Understanding_i$ assigns values from 4 to 1 to answers "strongly agree" to "strongly disagree" to the question "Do you understand this product assortment algorithm?" $Experience_i$ assigns values from 4 to 1 to answers "More than 10 years", "More than 5 years no greater than 10 years", "More than 1 years no greater than 5 years", and "No greater than 1 years" to the question "How many years did you work for this business?".