THE BEHAVIOURAL INSIGHTS TEAM

Alfred vending machines trial

Brief Analysis Report

Executive Summary

This note summarises the analysis undertaken for the Alfred Health vending machines trial. We report the results of analysis conducted at three levels – transaction level, machine level and product level.

At a transaction level, the trial found:

- That there was a 6.6 percentage point drop in sales for "red" drinks,
- A 3.8 percentage point increase for "amber" drinks
- 2.8 percentage point increase for "green" drinks.
- All of these effects are significant at the conventional level.

At the vending machine level, we see similar results to the estimates at the transaction level, specifically:

- A significant reduction of 8.3% in "red" drinks,
- A significant increase of 5.1% of "amber" drinks
- An increase of 3.2% in "green" drinks.

We also include analyses at a product level, which have been previously reported. At the product level, we estimate:

- A similar impact on "red" drinks, a reduction of 6.9%, but a large reduction in "amber" (32.6%) and a larger increase in "green" drinks (19%).

Due to a coding error, the estimates at a product level differ to those reported elsewhere. We would recommend reporting the analyses at the transaction or vending machine level.

Introduction

Background

This document reports the results from the evaluation of the Alfred Health vending machines trial. In this trial, we tested the impact of a 20% price increase on the consumption of high-sugar drinks, which were designated as "red drinks" on the basis of their nutrition rating. These ratings were assigned by Alfred Health using their standard processes. We randomly allocated half of the vending machines within Alfred Health Hospitals to have higher prices for their red drinks, while all other prices remained unchanged.

This trial was conducted as part of a wider body of work exploring interventions aiding individuals to make healthier decisions as part of the VicHealth Leading Thinkers Program. The trial ran for 6 months between August 2015 and January 2016, across 16 vending machines.

Data

Description

The original data consists of 864 rows, with each row representing one unique combination of vending machine and product. Each row then has 43 variables, 37 of which represent the number of sales of that product in that vending machine in a given month, from January 2013 through to January 2016. We have classified each product in to one of three categories. "Green" drinks, which are water. "Amber" drinks, which are sugar-free substitutes for high-sugar beverages (such as Pepsi Max) and "Red" drinks, which are high sugar drinks. This includes juice, sports drinks and traditional sodas. A chart of the total sales of each type of drink in control and treatment is provided below.

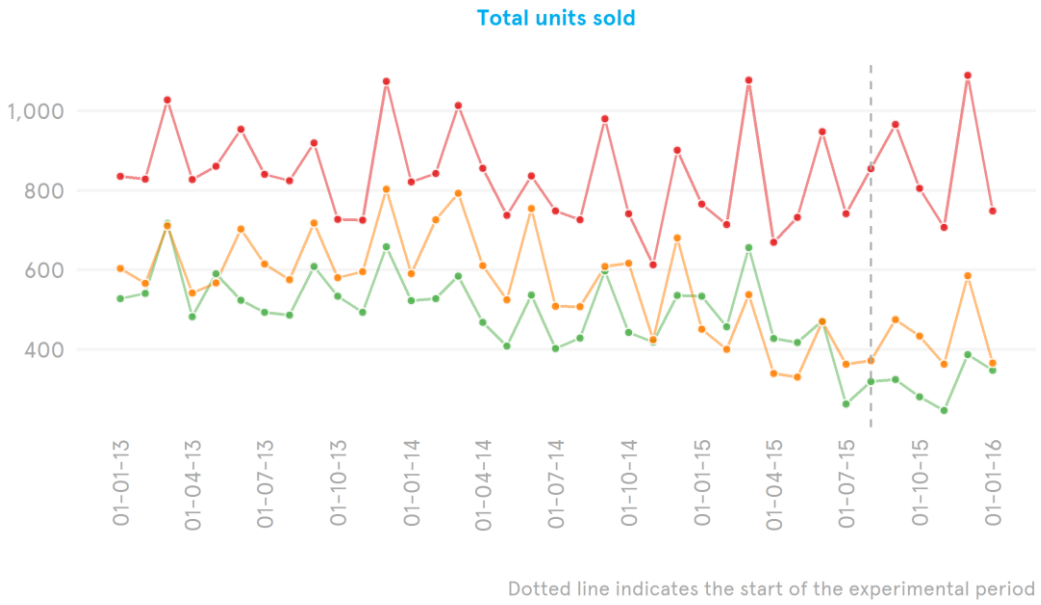

When we convert the data from wide to long format, where each row is then an observation of one good in one vending machine in a month, we have 31,968 observations.

We then select from the data just those observations that correspond to the vending machines and goods that we are interested in for this analysis. This reduces our sample size to 5,365 observations.

We then remove all observations where there are 0 sales of a given product in a given month, leaving us with a final sample size of 3,835 at the product level. Here, each row corresponds to one good at one vending machine in one month in which there were possible sales.

There are also two other levels at which we could conceptualise the data. One approach would be to think of each sale as an observation. When we transform the data to reflect this, we have 69,115 observations. Similarly, if we use each vending machine/month as a single observation, then we have 571 observations (the number of vending machines times the number of months for each vending machine was active).

As these are transformations of the data, the decision to analyse the data at a different level is implicitly a choice of weighting of information in the data-set. For transparency, we will provide estimates at all three levels of analysis.

Summary statistics

Below are summary statistics at the vending machine level.

	Total sales	Products	% Red	% Amber	% Green
Alfred Hospital Fairfield House	443	8	58%	30%	12%
Caulfield Medical Centre Cante	915	9	52%	35%	12%
Alfred Hospital Lvl 2 Sth Bloc	1,013	6	47%	53%	0%
Alfred Hospital Operating Thea	1,502	9	49%	42%	8%
Alfred Hospital Ward F Lvl 1	2,013	9	39%	41%	20%
Alfred Hospital Lvl 5 Lifts	2,237	8	57%	22%	21%
Alfred Hospital Physc Ward Gr	3,487	9	32%	63%	5%
Alfred Hospital Lvl 7 Lifts	4,180	8	50%	29%	21%
Alfred Hospital Lvl 4 Lifts	4,236	9	47%	30%	23%
Caulfield Medical Centre Breez	4,370	9	52%	36%	13%
Alfred Hospital Lvl 3 Lifts	4,559	9	50%	27%	23%
Alfred Hospital Lvl 6	4,742	9	49%	31%	20%
Alfred Hospital Lvl 2 Lifts	5,967	8	51%	28%	21%
Alfred Hospital Gr Fl Breezewa	6,892	11	34%	31%	35%
Alfred Hospital Amrep Hallway	6,990	11	38%	25%	37%
Alfred Hospital Emergency	15,569	13	44%	19%	36%

It's clear that there are substantial differences between machines. For example, Alfred Hospital Emergency has more than double the number of sales compared to

the second largest vending machine. Also note that the composition of drinks sold in each machine varies significantly – the proportion of drinks that are "red" varies from 34% to 58% of drinks.

Balance checks

Below we present the same summary statistics by condition, using only observations from the experimental period:

	Total sales	Products	% Red	% Amber	% Green
Control	4599	10	56%	29%	15%
Treatment	5065	11	51%	24%	24%

As you can see from the summary statistics, there is a slight compositional difference in the drinks sold, with fewer "green" drinks being sold in the control group than in the treatment group. Due to the extremely large Alfred Hospital Emergency machine, this also means that the relative size of the machines differs between control and treatment. This can be seen in the chart below, where each bar represents a single vending machine in each condition.

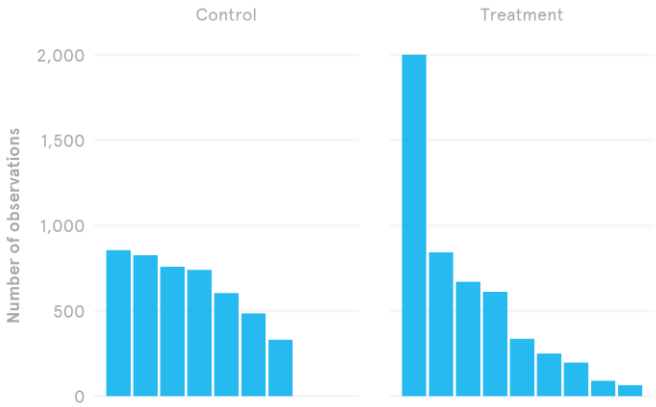

This poses some challenges for our analysis, as our randomisation across vending machines has resulted in non-trivial differences in balance at the product and transaction level.

Analysis

Primary analysis

Below, we present analogous analysis at the transaction, product and vending machine level. For each level of analysis, we present the percentage point change in sales for each type using an OLS regression, using indicator variables for each time period and fixed effects for each vending machine.

Transaction level

At the transaction level, we estimate three regressions of the form:

$$Y_{it} = \alpha + \beta_1 T_{it} + \beta_2 Vending\ Machine_{it} + \beta_3 Time_t + \varepsilon_{it}$$

Where Y_{it} is an indicator variable which is 1 when the transaction is of the type being estimated (i.e. 1 for Red for the Red estimate, and 0 otherwise). This means that we estimate the treatment effect on each type of drink in isolation. T_{it} is the treatment indicator, $Vending\ Machine_{it}$ is a vector of indicator variables for each vending machine, and $Time_t$ is a series of indicator variables for each time period.

The treatment indicator is coded as 0 for all observations before the experimental period, and for observations that happen within control vending machines during the experimental period, and 1 otherwise.

This means that β_1 estimates the average change in the sales of each type caused by the change in price, expressed as a percentage of sales.

The regression table below presents the results of these three regressions.

	Dependent variable: Transaction was of type Red/Amber/Green (YES or NO)		
	Red	Amber	Green
	(1)	(2)	(3)
Impact of price increase	-0.066***	0.038***	0.028**
	(0.011)	(0.010)	(0.009)
Controls for time	Yes	Yes	Yes
Controls for vending machine	Yes	Yes	Yes
Constant	0.362***	0.252***	0.387***
	(0.012)	(0.011)	(0.011)
Observations	69,115	69,115	69,115
Adjusted R^2	0.025	0.055	0.061
Note:			*p<0.05; **p<0.01; ***p<0.001

We estimate that the increase in prices decreased the consumption of "red" drinks by 6.6%, increased the consumption of "amber" drinks by 3.8% and increased the consumption of "green" drinks by 2.8%. These treatment effects are shown graphically below.

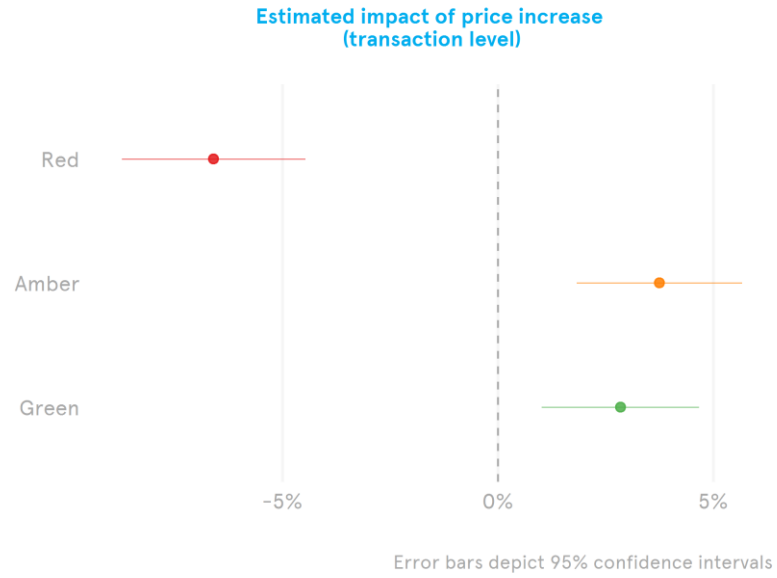

Product level

At the product level, we estimate three regressions of the form:

$$\ln(Y_{it}) = \alpha_0 Type_{it} + \beta_1 Type \times T_{it} + \beta_2 Vending\ Machine_{it} + \beta_3 Time_t + \varepsilon_{it}$$

Where Y_{it} gives the number of units sold for each product in each vending machine. Again, this means that we estimate the treatment effect on each type of drink in isolation. $Type_{it}$ is an indicator variable which is one if the product is of the type being estimated (i.e. 1 for Red for the Red estimate, and 0 otherwise) T_{it} is the treatment indicator, $Vending\ Machine_{it}$ is a vector of indicator variables for each vending machine, and $Time_t$ is a series of indicator variables for each time period.

The treatment indicator is coded as 0 for all observations before the experimental period, and for observations that happen within control vending machines during the experimental period, and 1 otherwise.

Again, this means that β_1 estimates the average change in the sales of each type caused by the change in price, expressed as a percentage of sales.

The regression table overleaf presents the results of these three regressions.

	Dependent variable: Transaction was of type Red/Amber/Green (YES or NO)		
	Red	Amber	Green
	(1)	(2)	(3)
Impact of price increase	-0.069	-0.326***	0.190
	(0.072)	(0.126)	(0.162)
Controls for time	Yes	Yes	Yes
Controls for vending machine	Yes	Yes	Yes
Observations	3,835	3,835	3,835
Adjusted R^2	0.916	0.911	0.889

Note: *p<0.05; **p<0.01; ***p<0.001

Also note that a transformation has been made to create percentages from the natural logs

At the product level, we estimate a similar impact on "red" drinks, a reduction of 6.9%, but now see a large reduction in "amber" (32.6%) and a larger increase in "green" drinks (19%). These treatment effects are shown graphically below.

The combination of high standard errors on the coefficients of interest, markedly different coefficient estimates with unreasonably high R^2 values suggests that there may be significant issues with the analysis at this level. This is likely to be driven by differences in balance and weighting of observations, exacerbated by the log transformation of our outcome variable.

Vending machine level

Finally at the vending machine level, which can also be considered the vending machine by month level, we estimate three regressions of the form:

$$Y_{it} = \alpha + \beta_1 T_{it} + \beta_2 Vending\ Machine_{it} + \beta_3 Time_t + \varepsilon_{it}$$

Where Y_{it} is the proportion of sales in each vending machine of the corresponding type. This means that we estimate the treatment effect on each type of drink in isolation. T_{it} is the treatment indicator, $Vending\ Machine_{it}$ is a vector of indicator variables for each vending machine, and $Time_t$ is a series of indicator variables for each time period.

The treatment indicator is coded as 0 for all observations before the experimental period, and for observations that happen within control vending machines during the experimental period, and 1 otherwise. Again, this means that β_1 estimates the average change in the sales of each type caused by the change in price, expressed as a percentage of sales. The regression table below presents the results of these three regressions at the vending machine level.

	Dependent variable: Transaction was of type Red/Amber/Green (YES or NO)		
	Red	Amber	Green
	(1)	(2)	(3)
Impact of price increase	-0.083***	0.051*	0.032
	(0.024)	(0.023)	(0.018)
Controls for time	Yes	Yes	Yes
Observations	571	571	571
Adjusted R^2	0.174	0.124	0.123

Note: *p<0.05; **p<0.01; ***p<0.001

At the vending machine level, we see similar results to the estimates at the transaction level, with an estimated reduction of 8.3% in "red" drinks, an increase of 5.1% of "amber" drinks and an increase of 3.2% in "green" drinks. These results are shown graphically below.

These estimates are qualitatively very similar those from the transaction level analysis, but again differ from the product level analysis where we saw a strong negative effect on "amber" drinks. Our estimates are less precise than analysis at the transaction level, which is expected due to the sharp drop-off in observations.

Summary

A summary of the estimates is provided below. In a table, and they are also presented graphically to give a sense of the difference in precision.

	Red	Amber	Green
Transaction	-6.6%	3.8%	2.8%
Product	-6.9%	-32.6%	19.0%
Vending Machine	-8.3%	5.1%	3.2%

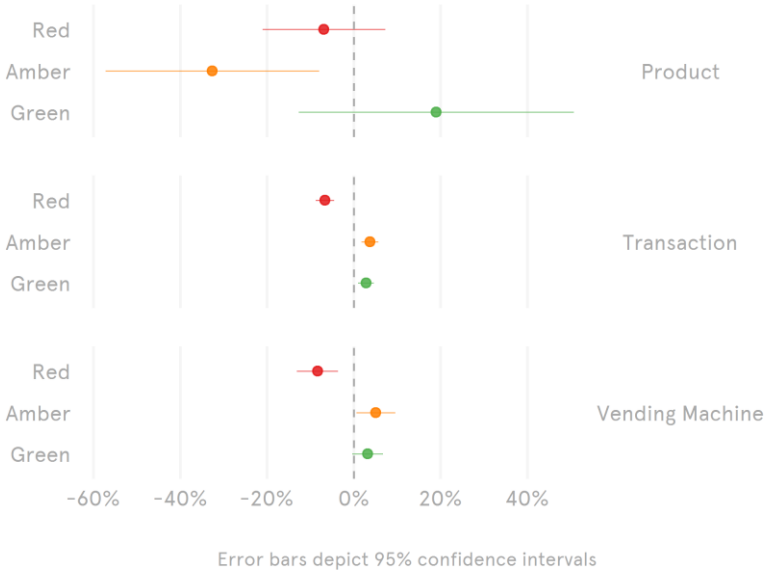

Our estimate of the impact on red drinks stays relatively constant, between 6%-9%, while the Amber and Green drinks fluctuate at the product level. As our transaction level estimates are the most conservative, we recommend using these results as definitive.

Appendix

Allocation of drinks

Below are the drinks that were observed in the trial, and the 'type' that they were allocated to, based on their sugar content.

Product		Type	Total sold
Frantelle	600Ml/24	Green	8,330
Cool Ridge Water	600Ml/24	Green	9,318
Pepsi Max	375Ml/24	Amber	1,303
Pepsi Max	600Ml/24	Amber	19,097
Mountain Dew	375Ml/24	Red	17
Gatorade Blbolt	600Ml/12	Red	21
Schweppes Solo	375Ml/24	Red	55
Gatorade Tropical	600Ml/12	Red	58
Schweppes Sunkist	375Ml/24	Red	94
Mineral Or & Mango	375Ml/24	Red	197
Sch Lemonade S/F	375Ml/24	Red	593
Schweppes Lemonade	375Ml/24	Red	903
Pepsi	600Ml/24	Red	2,157
Lipt Ice Tea Peach	500Ml/24	Red	3,252
Gatorade Orange	600Ml/12	Red	3,809
Svalley Orange	350Ml/18	Red	5,077
Svalley Apple	350Ml/18	Red	5,266
Lemonade	600Ml/24	Red	9,568