APPENDICES

EC.1. Additional Tables

To check our observations statistically, we use a left-censored panel regression with subject-specific random effects. Note that subjects were randomized across treatments, hence the standard assumption of orthogonality between random effects and covariates is satisfied. Left censoring is necessary because more than half of the observations in the dataset achieve an optimality gap equal to 0%, i.e., the subject chose the optimal assortment. Table EC.1 reports the regression coefficients and their significance, with standard errors provided in brackets. We note that both $\log\sigma_\mu$ and $\log\sigma_\nu$ are statistically significant and close in magnitude. This supports the appropriateness of a random effects model, as both individual heterogeneity and idiosyncratic variability substantially contribute to the variation in the dependent variable.

For clarity of exposition, we use NS as the base level for the decision support factor and HOV as the base level for the optimal variety factor. We observe that within the HOV condition, the optimality gap under the PB condition is approximately 4.55 percentage points larger than under the NS condition, with this difference being statistically significant, indicating that providing subjects with probabilities of buying worsens their performance compared to the NS condition, where no decision support is provided. However, within the LOV scenario, the PB information is associated with a significant reduction of the optimality gap by approximately 12.74 percentage points (calculated as $4.547 - 17.291 = -12.744$). Thus, the effect of PB is significant in both conditions, but in opposite directions: it is detrimental within the HOV condition but beneficial within LOV. We also note that the Period coefficient is negative and significant, indicating a learning effect: subjects make better assortment decisions over time, as reflected in a narrowing of the optimality gap.

Table EC.1 Experiment 1 Optimality Gap: Censored Panel Regressions with Random Effects (%)

	Regression 1	*Regression 2*
Intercept	16.758*** (1.214)	16.514*** (1.232)
Period	−0.569*** (0.027)	−0.550*** (0.031)
LOV	−0.533 (1.454)	−0.003 (1.734)
PB	4.547*** (1.411)	4.533*** (1.413)
EP	−6.549*** (1.420)	−6.577*** (1.421)
Period×LOV		−0.044 (0.060)
PB×LOV	−17.291*** (1.906)	−17.275*** (1.922)
EP×LOV	−1.420 (2.196)	−1.373 (2.198)
logSigmaMu	2.672*** (0.022)	2.672*** (0.022)
logSigmaNu	2.684*** (0.008)	2.685*** (0.008)
Total Observations	2,200	2,200
Left-Censored	1,150	1,150
Log Likelihood	−4,902.981	−4,902.981
Akaike Inf. Crit.	9,825.962	9,825.962
Bayesian Inf. Crit.	9,882.924	9,882.924

Note: *p<0.1; **p<0.05; ***p<0.01
LOV=Low Optimal Variety, PB=Probabilities of Buying, EP=Expected Profit

In Table EC.2, we report the results of panel regressions with random subject-specific effects, estimated separately for the BCS and SCS conditions. We separate the conditions because, as shown in Figure 6, the interaction effects between decision support, default assortment, choice set size, and learning complicate interpretation in a combined model.

We find that starting with all products prechecked (Pre) is associated with significantly worse performance in the BCS condition, increasing the optimality gap by 54.6 percentage points on average. In contrast, Pre significantly improves performance in the SCS condition, reducing the optimality gap by 4.44 percentage points. This is as expected: Pre may exacerbate mean-anchoring bias in large choice sets (BCS), while mitigating it in small choice sets (SCS).

Interestingly, the learning effect appears significant only for the BCS condition.

Table EC.2 Experiment 2 Optimality Gap (%) by Choice Set Size: Censored Panel Regressions with Random Effects

	BCS	SCS
Intercept	99.740 (4.464)	22.720*** (0.768)
Period	−3.274*** (0.516)	−0.038 (0.073)
PRE	54.619*** (3.178)	−4.442*** (0.842)
PB’	0.682 (6.611)	−2.800*** (0.993)
logSigmaMu	4.783*** (0.008)	2.342*** (0.030)
logSigmaNu	4.644*** (0.004)	2.661*** (0.006)
Total Obs.	2,552	2,552
Left-Censored	275	183
Log Likelihood	−14,128.220	−9,990.865
Akaike Inf. Crit.	28,270.440	19,993.730
Bayesian Inf. Crit.	28,311.360	20,028.800

Note: *p<0.1; **p<0.05; ***p<0.01. The first 3 periods dropped for outliers. BCS=Big Choice Set, SCS=Small Choice Set

Table EC.3 Experiment 2 SCS Optimality Gap (%). Outlier periods 1,2, 3 and 16, 17 are excluded

Intercept	23.471*** (0.824)
Period	−0.189** (0.078)
PRE	−5.051*** (0.983)
PB’	−0.752 (1.003)
logSigmaMu	2.357*** (0.035)
logSigmaNu	2.656*** (0.007)
Observations	2,320
Left-Censored	163
Log Likelihood	−9,097.105
Akaike Inf. Crit.	18,206.210
Bayesian Inf. Crit.	18,240.710

Note: *p<0.1; **p<0.05; ***p<0.01
SCS=Small Choice Set, Pre=Pre-checked, PB=Probabilities of Buying

Finally, the effect of PB' information is statistically significant only in the SCS condition, where it reduces the optimality gap by 2.8 percentage points on average. However, Table EC.3 shows that in the SCS condition, once the first two periods following the introduction of PB' information (periods 16 and 17) are excluded, the effect of PB' becomes statistically insignificant. This supports our conjecture that subjects may have found the PB' information confusing or cognitively burdensome. Although the assortments they selected with the aid of PB' were, on average, more profitable, many subjects may have been uncertain about whether they were using the information correctly and ultimately reverted to their prior strategies.

Table EC.4 Experiment 2 ANOVA on average assorment size per subject, periods 4 to 15

Factor	Df	Sum Sq	Mean Sq	F value	p-value
Default Assortment	1	237.9	237.87	44.97	$< 0.0001^{***}$
Cons. Set	1	296.1	296.07	55.97	$< 0.0001^{***}$
Default Assortment × Cons. Set	1	75.8	75.77	14.32	0.0002^{***}
Residuals	228	1206.0	5.29		

Note: Significance codes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

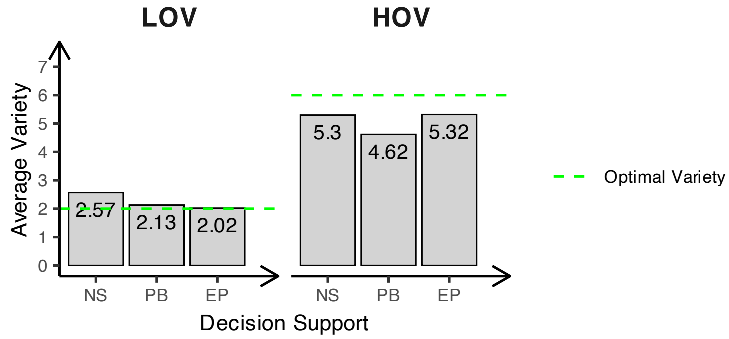

Figure EC.1 **Mean of subject-level assortment size means in Experiment 1 (popular assortments only)**

Table EC.5 **Experiment 2 Average Assortment Size: Tukey Post Hoc Comparison of Means with 95% family-wise confidence level.**

Comparison	Difference	Lower CI	Upper CI	p-value
Default Assortment				
PRE - UN	2.025	1.430	2.620	<0.001
Choice Set				
BCS - SCS	2.259	1.664	2.854	<0.001
Default Assortment × Choice Set				
PRE×SCS - UN×SCS	0.882	−0.223	1.987	0.168
UN×BCS - UN×SCS	1.116	0.011	2.222	0.047
PRE×BCS - UN×SCS	4.284	3.179	5.390	<0.001
UN×BCS - PRE×SCS	0.234	−0.871	1.340	0.947
PRE×BCS - PRE×SCS	3.402	2.297	4.508	<0.001
PRE×BCS - UN×BCS	3.168	2.063	4.273	<0.001

Note: p-values adjusted for multiple comparisons.
PRE = Pre-checked, UN = Unchecked, BCS = Big Choice Set, SCS = Small Choice Set.
Only periods 4 to 15 are included

In Table EC.4, we present the results of an ANOVA with the average assortment size per subject in *Experiment 2* as the dependent variable, and the corresponding post-hoc mean comparisons are shown in Table EC.5. To isolate the effects of default assortment and choice set size (i.e., relative optimal variety) from potential confounding due to the introduction of PB' information, we restrict the analysis to periods 4 through 15—after the initial learning phase and before PB' was introduced. (We note that extending the analysis to include periods 16 through 25 yields similar results.)

These findings confirm that assortment sizes differ significantly by choice set size, supporting the presence of a relative optimal variety effect: assortments are, on average, larger in the BCS condition, where the choice set size is large relative to the optimal variety, than in the SCS condition. Additionally, assortments tend to be larger when subjects begin each round with all products prechecked (Pre) compared to when they start with an empty selection (Un).

EC.2. Popular Sets Analysis

In this section, we verify the robustness of our findings by repeating the analyses for *Experiment 1* and *Experiment 2* while excluding the instances where subjects choose non-popular assortments. Overall, the results remain directionally consistent with those reported in the main body of the paper (where we include non-popular assortments), although some effects lose statistical significance due to the reduced number of observations.

Figures EC.1, EC.2, EC.3, EC.4, EC.5, and EC.6 correspond to Figures 3, 4, 7, 9, 6, and 8 in the main body of the paper respectively. We note that, after dropping non-popular assortments, the patterns in optimality gaps and assortment sizes remain consistent with those reported in the main text.

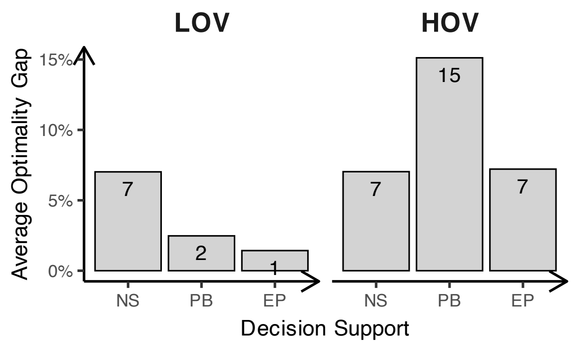

Figure EC.2 Mean of subject-level means of optimality gap in Experiment 1 (popular assortments only)

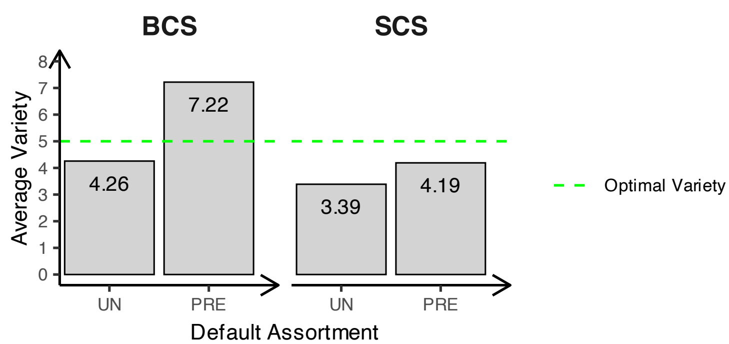

Figure EC.3 Mean of subject-level assortment size means in Experiment 2, periods 4-24 (popular assortments only)

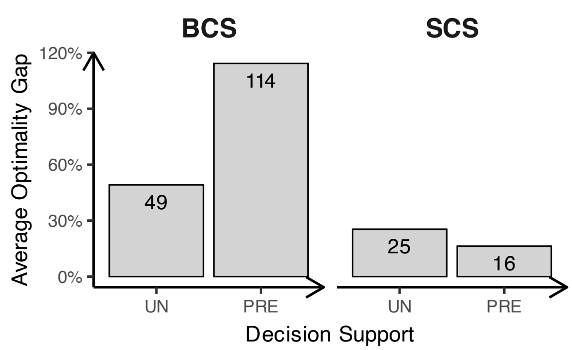

Figure EC.4 Mean of subject-level optimality gap means in Experiment 2 (popular assortments only)

Next, we statistically confirm our main result – a mean-anchoring effect for assortment sizes, – by repeating the analyses reported in Tables 2 and 4. When excluding non-popular assortments, the directional asymmetry tests can be implemented in two reasonable ways: one may compare the proportions of *too-small popular* assortment across treatments – i.e., $\hat{p}(\text{too small \& popular} \mid \text{HOV})$ versus $\hat{p}(\text{too small \& popular} \mid \text{LOV})$ – or, alternatively, compare the proportions $\hat{p}(\text{too small} \mid \text{popular, HOV})$ versus $\hat{p}(\text{too small} \mid \text{popular, LOV})$. We perform the tests both ways (reported in Tables EC.6 and EC.7 for *Experiment 1* and Tables EC.8 and EC.9 for *Experiment 2*) and find that the comparisons are directionally identical for both conditioning schemes, reflecting the clarity of the underlying mean-anchoring deviation pattern.

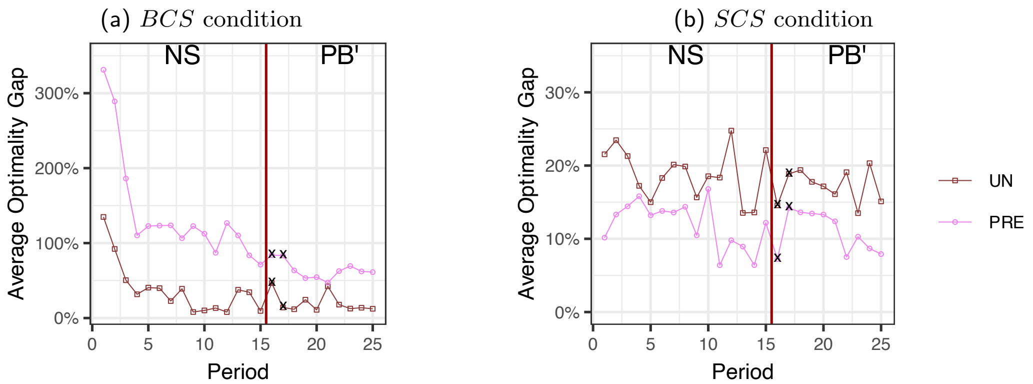

Figure EC.5 Overall subjects performance over time in Experiment 2 (popular assortments only)

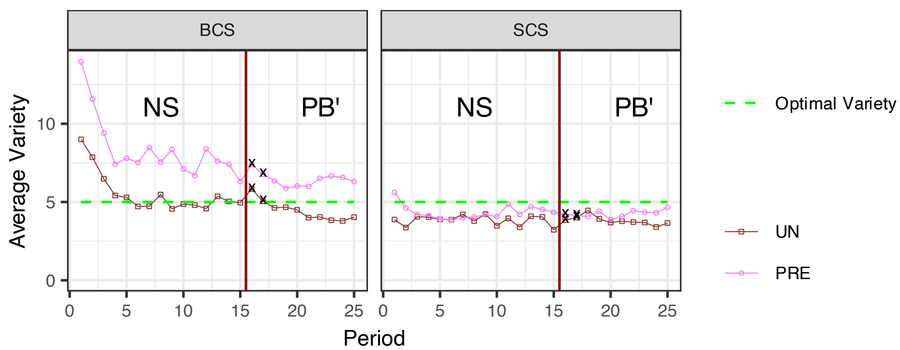

Figure EC.6 Average assortment sizes over time in Experiment 2 (popular assortments only)

Table EC.6 Average number of too small/large assortments per subject in Experiment 1 (popular assortments only)

	too small			*too large*		
	LOV	*HOV*	p-value	*LOV*	*HOV*	p-value
NS	1.08	8.50	<0.001***	7.25	2.83	0.0517*
PB	1.14	13.33	<0.001***	3.24	2.80	0.016**
EP	0.83	6.56	0.002***	1.50	2.31	0.895

Note *p<0.1; **p<0.05; ***p<0.01; LOV=Low Optimal Variety, HOV=High Optimal Variety
NS=No Support, PB=Probabilities of Buying, EP=Expected Profit

Table EC.7 Average % of too small/large assortments per subject in Experiment 1 (popular assortments only)

	too small			*too large*		
	LOV	*HOV*	p-value	*LOV*	*HOV*	p-value
NS	5.1%	40.8%	<0.001***	33.9%	13.6%	0.078*
PB	5.3%	61.5%	<0.001***	14.9%	12.9%	0.011**
EP	4.0%	29.5%	0.007***	7.3%	10.4%	0.895

Note *p<0.1; **p<0.05; ***p<0.01; LOV=Low Optimal Variety, HOV=High Optimal Variety
NS=No Support, PB=Probabilities of Buying, EP=Expected Profit

Table EC.8 Average number too small/too large assortments per subject in Experiment 2 (popular assortments only)

	too small			*too large*		
	BCS	SCS	p-value	BCS	SCS	p-value
Un	5.64	8.55	0.01^{**}	3.19	2.45	0.17
Pre	5.36	6.50	0.22	9.31	2.86	$<0.001^{***}$
p-value	0.30	0.03^{*}		$<0.001^{***}$	0.02^{**}	

Note *p<0.1; **p<0.05; ***p<0.01,

BCS=Big Choice Set, SCS=Small Choice Set, Un=Unchecked, Pre=Pre-checked

Table EC.9 Average % of too small/large assortments per subject in Experiment 2 (popular assortments only)

	too small			*too large*		
	BCS	SCS	p-value	BCS	SCS	p-value
Un	50.8%	68.8%	0.02^{**}	28.7%	19.7%	0.09^{*}
Pre	30.6%	57.1%	$<0.001^{***}$	53.1%	25.2%	$<0.001^{***}$
p-value	$<0.001^{***}$	$<0.001^{***}$		$<0.001^{***}$	0.007^{**}	

Note *p<0.1; **p<0.05; ***p<0.01,

BCS=Big Choice Set, SCS=Small Choice Set, Un=Unchecked, Pre=Pre-checked

EC.3. Analysis of the Clickstream Data

As mentioned previously, in addition to the choices submitted by the subjects, we collected their clickstream data. Specifically, in *Experiment 1*, we recorded the clicks on the check/uncheck boxes to select and unselect products in the assortment and the hitting of the "Submit" button, along with the time stamps of those clicks. We used JavaScript to collect the clickstream data, which sometimes gave inaccurate results due to technical issues with the user's browser which were beyond our control. We verified the consistency of the collected clickstream data with the final decision each subject submitted to the server and discarded the observations that were not consistent with the submitted selections, leading to a total of 22 discarded observations out of a total of 2200 and affecting 7 out of 88 subjects. In what follows, when reporting period averages, we include all subjects; when reporting cumulative values (e.g., the number of different assortments seen in the whole experiment), we include only the 81 subjects for whom we have complete entries for all 25 periods.

The purpose of our analysis of the subjects' clicking behavior is to get some insights into their decision making process. To this end, we introduce the following definitions. In *Experiment 1*, we say that an assortment is *seen* by a subject if at some point during a period the corresponding product boxes were checked simultaneously. We say that an assortment is *chosen* by a subject if it was submitted in a given period. The number of clicks made by a subject in a given period includes the clicks that were made to check and uncheck the boxes. For example, if in a given period, starting from all product boxes unchecked, a subject checks the product A box, then checks the product D box, then unchecks the product A box, then checks the product A box again and clicks "Submit", the sequence of assortments seen is {A}, {A, D}, {D} and {A,D} and the chosen assortment is $\{A, D\}$. In our accounting, this would correspond to a total of four clicks and three different assortments seen. Note that we do not count empty assortments, because subjects always start each period with all boxes being unchecked, so the empty assortment is seen in each period.

Figure EC.7 shows averages for the number of different assortments seen by the subjects in each period. We observe that subjects see more assortments in the first few periods of the experiment, which is consistent with an early exploration behavior. Recall that all the parameters remain constant throughout our experiment, so that the optimal strategy was to select the same (expected profit-maximizing) assortment in each of the 25 periods.

Note that the subjects in the *HOV* condition on average see more different assortments than the subjects in the *LOV* condition, but this does not necessarily imply that they are more inclined to explore: in the *HOV* condition a subject who chooses the optimal assortment (which contains six products) in the *HOV* condition,

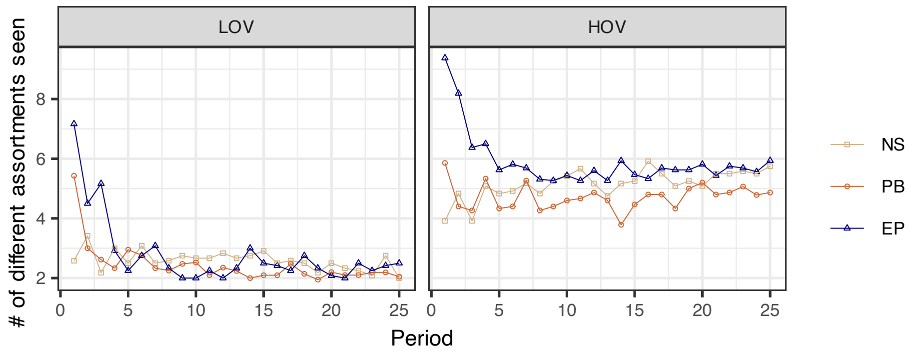

Figure EC.7 **Experiment 1. Average number of different assortments seen per period.**

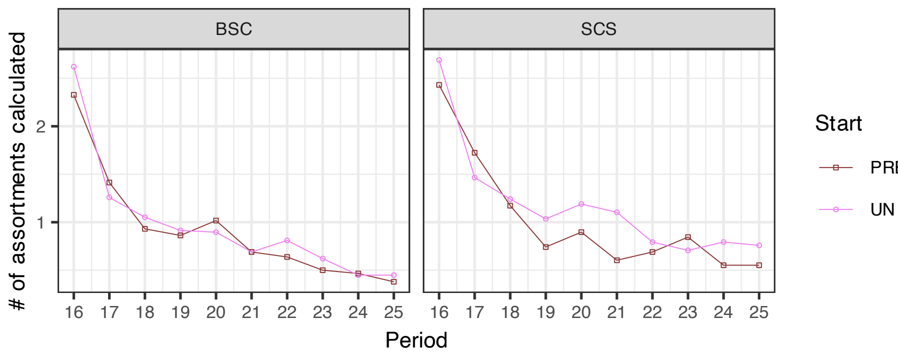

Figure EC.8 **Experiment 2. Average number of assortments calculated per period.**

will see at least six different assortments in that period; in comparison a subject in the *LOV* condition may see as little as two different assortments when choosing the optimal assortment.

Table EC.10 **Experiment 1. Mean number of different assortments seen and chosen**

	Seen	Chosen	Subjects
HOV	28.10	6.05	40
NS×HOV	28.91	7.09	11
PB×HOV	22.64**	6.43	14
EP×HOV	32.6	5	15
LOV	13.66	5.34	41
NS×LOV	13.9	6.73	11
PB×LOV	11.72	5.39	18
EP×LOV	16.33	4*	12
Grand Total	20.79	5.70	81

Table EC.11 **Experiment 1. Mean number of clicks in certain periods**

	1st	5th	25th
HOV	9.31	5.37	5.65
NS×HOV	4.00	4.83	5.75
PB×HOV	7.00**	4.47	4.93*
EP×HOV	15.31***	6.63**	6.25
LOV	9.53	3.11	2.29
NS×LOV	2.67	2.50	2.00
PB×LOV	9.38***	3.43	2.05
EP×LOV	16.67***	3.17	3.00
Grand Total	9.43	4.21	3.93

Note: *p<0.1; **p<0.05; ***p<0.01. *Mann-Whitney U test for equality in ranks with NS condition*

To make judgements about the subjects' exploration behavior we compare the number of assortments seen in the *EP* and *PB* conditions to the ones seen in the *NS* condition, where subjects did not get any information from clicking. In Table EC.10 we see that subjects in the *EP* condition *see* more different assortments but *choose* fewer different assortments than in the *NS* condition, which suggests that they indeed take advantage of the information provided in the *EP* condition. Mann-Whitney U tests, however, do not indicate significance, except for one case (chosen assortments within the *LOV* condition). Curiously, it appears that subjects in the *PB* condition see significantly fewer assortments than subjects in the *NS* condition. This finding further supports our previous observation that providing subjects with purchase probabilities hinders their performance in the *HOV* condition.

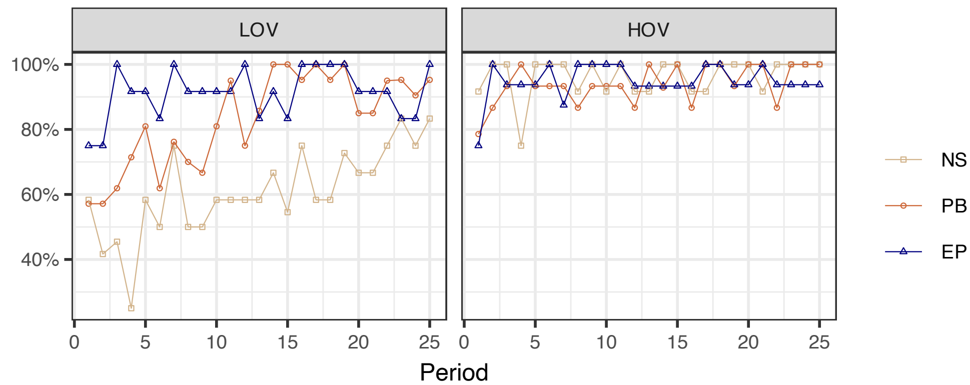

Figure EC.9 Experiment 1. Percentage of subjects who picked the best assortment among the ones seen

In Figure EC.9 we represent the percentage of subjects who, in each period, chose the assortment which generated the highest expected profit of all the ones they saw in that period. We are specifically interested in the *EP* condition, where the expected profit values were clearly displayed for each assortment seen by the subject. When we look at the values for the first period, we observe that not all the subjects in that condition chose the assortment with the maximum expected profit value among the ones they saw: only 75% of them did so.

We provide four possible explanations for this surprising number. First, it is possible that some subjects think that choosing the assortment with the highest expected profit is not the optimal strategy due to risk aversion. Second, subjects may have problems remembering the expected profits for all the assortments they have seen and comparing these numbers effectively. Third, despite our best effort to signal otherwise, subjects could suspect some kind of deception in our experiment. (Use of deception is discouraged in the scholarly field of behavioral operations management (Katok 2011); however, Amazon mTurk and Prolific participants are likely to have previous exposure to studies in other fields, for example, psychology, where ethical guidelines allow deception under certain conditions (American Psychological Association 2017).) Finally, it is possible that, even after successfully passing the instructions quiz, some subjects had problems understanding the concepts of "average profit" (recall that this is how we refer to expected profit in the instructions), random demand and period independence or struggled with getting used to the experiment interface. We see that in later periods almost all subjects in the *EP* condition chose the assortment with the maximum expected profit that they see in that period. Therefore, it seems that even though subjects may initially have reservations trusting or understanding the decision support information, their wariness or confusion diminishes over time.

In *Experiment 2*, starting in Period 16, subjects were given the option to click on a "Calculate" button in order to display the "chance of buying" for each product in the currently selected assortment. In this case, we say that the assortment was *checked* by the subject. Figure EC.8 shows that, even in Period 16, when this option was first offered to them, the subjects checked on average only about two assortments before submitting one, and most of them stopped clicking on the "Calculate" button after only two periods.

In Figure EC.10, we represent the percentage of subjects who, in Periods 16-25 chose the assortment which generated the highest expected profit of all the ones they checked in that period. We see that, here again, many subjects do not choose the highest expected profit assortment they saw the "chances of buying" for, which further confirms that this information may be confusing to subjects.

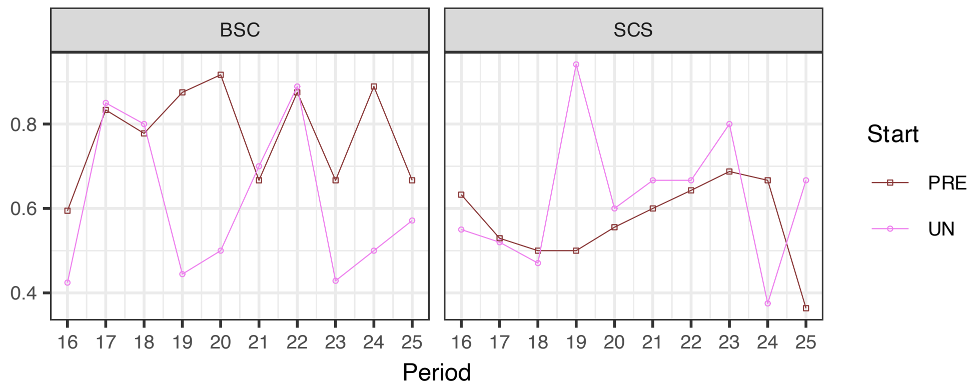

Figure EC.10 Experiment 2. Percentage of subjects who picked the best assortment among the ones they checked (only includes subjects who checked more than one assortment)

EC.4. Exit survey analysis

At the end of *Experiment 2* the subjects were asked to fill out an exit survey to elicit some further information about their strategy and perception of the task they had just completed, as well as to gather some demographical information. Summary results for the demographical information are presented in Table EC.12.

Gender		Age		Education		Employment status	
Male	146	18-24	126	High school	80	Employed	122
Female	71	25-34	47	Associate	17	Student	77
Other	5	35-44	34	Bachelor	73	Unemployed	20
		45-54	15	Master	48		
				Doctorate	2		
				other	2		

Table EC.12 Demographic information

The exit survey included an *attention check question* (Question 4) meant to ensure the subjects were indeed paying attention and not just clicking on random answers for the sake of completing the survey as quickly as possible. Out of the 230 participants in *Experiment 2*, 222 (97%) passed the attention check question. Those who did not were excluded from further analysis. Although there was no reward associated with participation in the post-experiment exit survey (i.e., the subjects were not paid for their time spent filling it), 130 participants (57%) took time to answer the open-ended essay question inviting them to share "insights regarding how [they] approached the game and the strategy [they] adopted" (Question 10) and 30 participants provided additional information regarding the factors they considered in their decision process (Question 2).

When asked which factors they considered when selecting their assortment in each period, 79.3% of the subjects mentioned the popularity levels of the products, as shown in Figure EC.11.

When asked whether the best assortment was the same in every period or varied, 40.09% of the subjects correctly answered it was the same and 50.45% said it varied. Regarding optimal variety, 22.97% of the subjects correctly identified the optimal number of products as five but interestingly, many more underestimated this number (54.95% reported a number between 0 and 4) than overestimated this number (18.02% reported a number between 6 and 20).

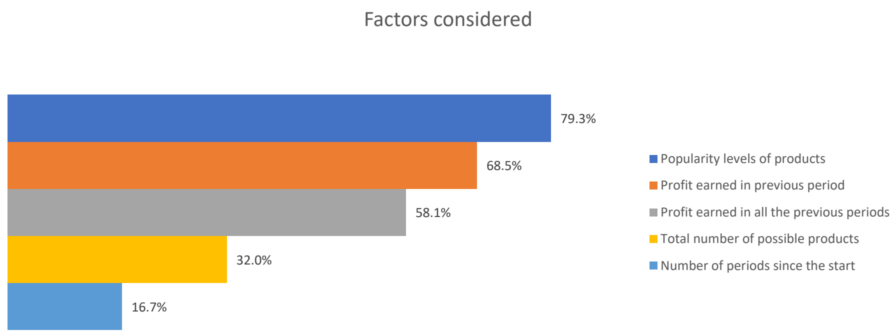

Figure EC.11 **Exit survey: Factors considered by subjects**

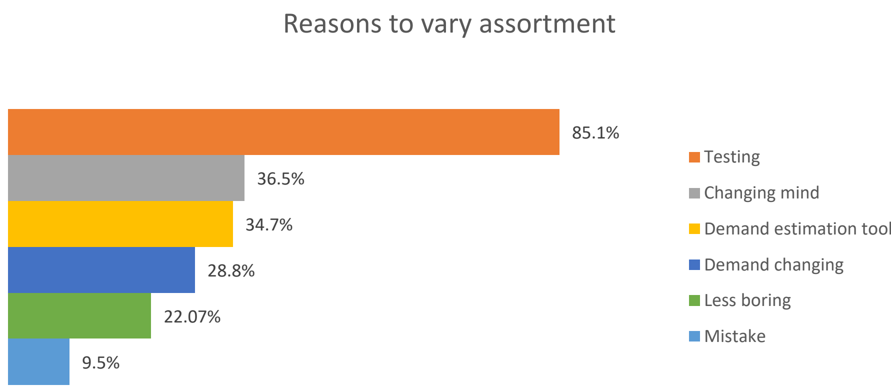

Figure EC.12 **Exit survey: Reasons to vary the assortments as reported by subjects**

Regarding the reasons why the subjects varied their assortment over time, the most commonly cited reason, as shown in Table EC.12 was for testing purposes (85.1%) followed by the subjects changing their mind (36.5%).

Interestingly, 34.7% of the subjects mentioned that the demand estimation tool which was introduced in Period 16 was one of the reasons why they varied their assortment. When specifically asked about the impact of the PB information, 37.8% of the subjects said it encouraged them to vary the assortment more from period to period, 36% said it made them realize they were selecting too many products, as seen in Figure EC.13. Only 17.1% said it made them realize they were selecting too few products. This seems to suggest that the PB information may be contributing to the censoring of demand information which leads to subjects choosing too small of assortments.

We also asked two questions with the aim of checking the understanding of subjects of the role of the popularity indices in the optimal assortment selection. To the hypothetical question "suppose one more product with popularity 1 (resp. 50), i.e., lower (resp. higher) than all the other ones, is added to the list of possible products, would your assortment strategy change?", the majority of subjects give the correct answer, specifically 59.9% replied that it would stay the same for the popularity index of 1 and 81.5% said it would change for the popularity index of 50. This provides further support for the assertion that subjects understood the importance of the popularity indices in the assortment decision.

Upon careful inspection of the subjects' open-ended answers regarding strategy, many reported following some kind of "trial-and-error" process, saying "*once I found the assortment that seemed to consistently result in the highest actual profit, I stuck with it*" or "*At first it was trial and error. Then I found an assortment*

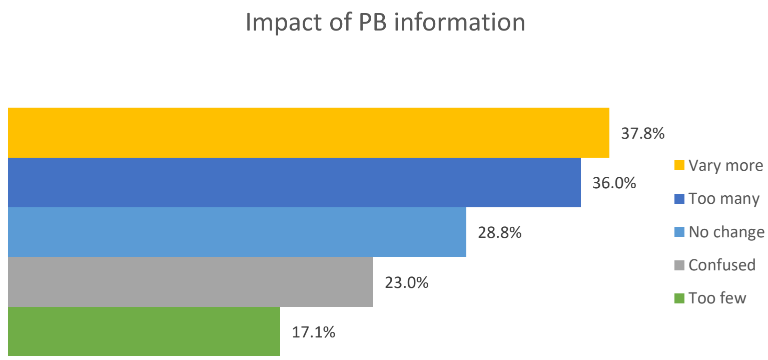

Figure EC.13 Exit survey: Impact of the PB information

that consistently gave me 33 so I stuck with that but also tried new ones ever now and again" or "*I mainly used the "trial and error" method before sticking with one assortment.*". Some subjects specifically mentioned starting with a large number of products then reducing variety over time :"*I started with almost a random number and realized it was way too many products, so I was losing money. Then, I lowered it and finally found a number or range (about 3-6) that consistently got me a lot of money.*" or "*I began the game by selecting a wide variety of products in order to see how they did on the market. Then I decreased the amount of products and tested every now and then to see which products and the amount of products would return the most profit*".

Some subject explicitly recognized focusing on popular sets "*My strategy was to pick a couple of products with the highest popularity index*" or "*The strategy I adopted was selecting high popularity index products that had a higher chance of being picked.*" or "*I tried to pick the highest popularity while also varying my choices a bit to see if the outcomes might change*" or " *The three least popular products did not seem to be very profitable, so I tried to stay away from those*".

Some comments point towards a fundamental mis-understanding of the random nature of demand and profit and to the misinterpretation of feedback by subjects as previously noted in newsvendor experiments (Bolton and Katok 2008, Schweitzer and Cachon 2000): "*Even if popularity was high, it did not mean selling it was going to be profitable*" or "*it felt like the computer was calculating at random so I chose at random*" or "*Towards the start of the experiment and the beginning of the second section, I did my best to calculate what the earnings might be. The actual EUs earned were always a little different than what I calculated, which may have been error on my part, or the percentages were more of a 'chance' than a fixed, for sure number. [...]* ".

Some subjects seem to be expressing loss aversion: "*I found a way that would never produce negative profits*" or "*I am a very safe person. Once I saw there was profit, I would tend to stay towards selling those same products regardless of the margin of profit. A profit range of 10-15 was acceptable to me. I really only changed strategies once I noticed I was in the negative. I am not a greedy person.*"

Some subjects in the Pre-checked treatment expressed frustration with having to un-check lots of products in each period: "*I wish there was a deselect all button or they were not all selected automatically*".

Some subjects spontaneously praised the probability calculations offered starting in Period 16: "*I liked the calculator*", and "*when the game allowed me to calculate the % it was much easier to form various assortments*".

Finally the following comments, although anecdotal, signaled to us that subjects tend to over-think the problem by assuming there is some "magical formula" to crack the game and also speaks to the tradeoff between exploration and exploitation: "*In the back of my head, I knew that there must be a better way to maximize profit than simply choosing the top 3-4 options, and thus, would experiment now and again with the less popular items; however, towards the end, I ended up sticking with the safer option of at least some guaranteed profit rather than possibly losing more by experimenting as often as did in the beginning. Though,* *the experimentation felt necessary in order to get a feel for the mechanics of the game.*" and "*I tended to stick with the more popular product, but switched things up occasionally to see whether there was some underlying mechanic that was being considered.*"

EC.5. Random Utility Model

A natural way to extend this work is to develop and fit a structural model for predicting the behavior of assortment planners. An intuitive way to do this would be to use a random utility model (RUM), where the nominal utility is based on the expected profit of the assortment, with added behavioral components, such as preferences for popular assortments, a mean-anchoring effect, and a distance to the default assortment.
If, like with the consumer choice model, we assume that the bias parameters are the same for all assortment planners, and the error terms follow Gumbel distribution, we would end up with an MNL choice model defined over the output of the consumer's MNL choice model.

For illustration, we may write the utility function as:

$$U_{it}(S) = \mathbb{E}[\pi(S)] + \gamma_1 I_{pop} - \gamma_2 \left(|S| - \frac{n}{2}\right)^2 - \gamma_3 \left(|S| - d\right)^2 + \nu_i + \zeta_{it},$$

where I_{pop} is a binary variable indicating whether the assortment S is popular, d is the size of the default assortment, ν_i is the subject-specific effect, and ζ_{it} is the idiosyncratic error. Recall that the expected profit is

$$\mathbb{E}[\pi(S)] = r\lambda \left(1 - \frac{v_0}{\sum_{j \in S} v_j}\right) - K \times |S|,$$

If we remove the behavioral terms (i.e., set all γ coefficients to zero), this model reduces to a *Quantal Response Equilibrium (QRE)* model, where assortment choices depend only on expected profit, with a "trembling hands" effect modeled by a bounded rationality parameter β:

$$P(S \in \mathcal{P}(N) \mid \beta) = \frac{e^{\pi(S)/\beta}}{\sum_{T \in \mathcal{P}(N)} e^{\pi(T)/\beta}},$$

This model provides a useful benchmark: under QRE, the most frequently chosen assortment should be the one with the highest expected profit, and learning should cause the distribution of choices to become more concentrated around that optimum. However, our data strongly suggest that this baseline model is incomplete. First, popular assortments are consistently chosen more often than non-popular assortments with similar expected profit, clearly indicating that $\gamma_1 > 0$. Second, we observe that the distribution of assortment sizes is not centered at the optimal value, but rather pulled toward the midpoint of the choice set. This implies that $\gamma_2 > 0$, consistent with the mean-anchoring bias documented earlier.

To illustrate this latter point, we compare our experimental data to the predictions of a QRE model constrained to popular assortments only. Due to the limited number of observations, we pool data from the *Un* and *Pre* treatments. For this restricted set $A_0, A_1, \ldots, A_n$ of popular assortments of size 0 through n, we compute

$$P(S = A_j \mid S \in \{A_0, A_1, \ldots, A_n\}, \beta) = \frac{e^{\pi(A_j)/\beta}}{\sum_{k=0}^{n} e^{\pi(A_k)/\beta}}$$

Figures EC.14a, EC.14b, and EC.14c show the observed distribution of choices and overlay the QRE predictions for three different experimental periods: the first, the last, and an intermediate period (period 10, selected arbitrarily). We visually select β to best match the spread of the data. As learning progresses, we observe some reduction in variance – as predicted by QRE – but the mode of the distribution does not align with the theoretical optimum. Instead, the entire distribution appears systematically shifted, with the shift toward the center of the choice set especially pronounced in the *SCS* treatment. Notably, the peak of the empirical distribution lies to the left of the QRE-predicted mode, suggesting the presence of anchoring or reference-dependent effects that are not captured by the standard QRE framework.

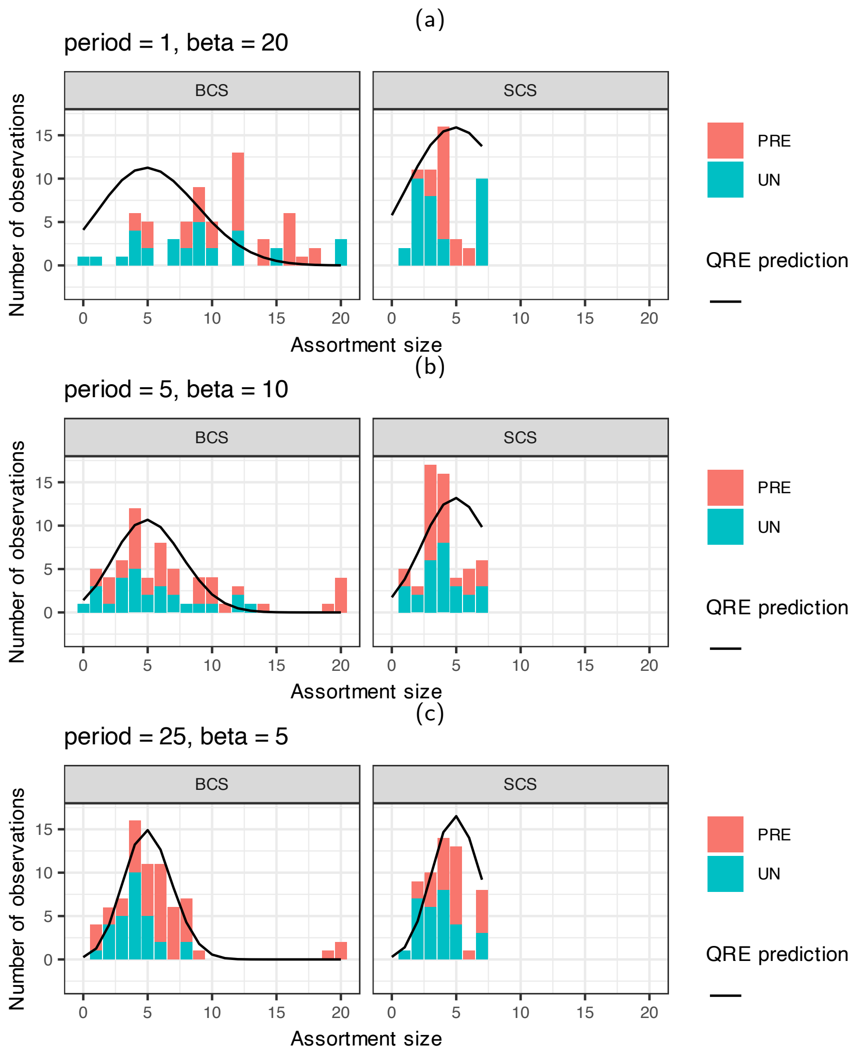

Figure EC.14 Experiment 2. Frequency of choosing popular assortments of each size in periods 1, 5 and 25.

We emphasize that our purpose here is not to estimate β or fit a structural model, but rather to visually demonstrate that the QRE model without behavioral adjustments (i.e., with $\gamma_2 = 0$) fails to match the data. This supports the inclusion of a nonzero pull-to-center term in future modeling efforts.

Estimation Challenges. While developing a behavioral RUM model for our setting is a promising research direction, estimating such models is highly challenging. Structural estimation requires maximizing a likelihood over the entire power set of N, i.e., 2^n possible assortment combinations. In our experiments, many of these combinations are never observed in the data, leading to identification and computational problems.

The literature offers several models for the empirical estimation and prediction of choices when a decision-maker selects a subset from a larger set. (e.g., Tran and Mai (2024), Sun et al. (2020)). However, these models assume additive utility, which does not hold here due to diminishing returns in $\mathbb{E}[\pi(S)]$. Specialized estimation techniques would be required, and currently available software is not well-suited to this context, as the expected profit is not additive in v-values. See Chapter 8 in Train (2009), as well as Chen et al. (2022) and Chen et al. (2012) for examples of the complexity involved in estimating even simpler RUMs. For these reasons, we do not attempt a structural estimation in this paper. However, we emphasize the need for future research to develop random utility models that are suitable for empirical settings involving non-additive utility structures.

References

Train KE (2009) Mixed logit. In: *Discrete Choice Methods with Simulation*, 2nd ed. (Cambridge University Press, New York).

Chen K-Y, Wang J, Lang Y (2022) Coping with digital extortion: An experimental study of benefit appeals and normative appeals. *Management Sci.* 68(7):5269–5286.

Chen Y, Su X, Zhao X (2012) Modeling bounded rationality in capacity allocation games with the quantal response equilibrium. *Management Sci.* 58(10):1952–1962.

EC.6. Subjects Screen

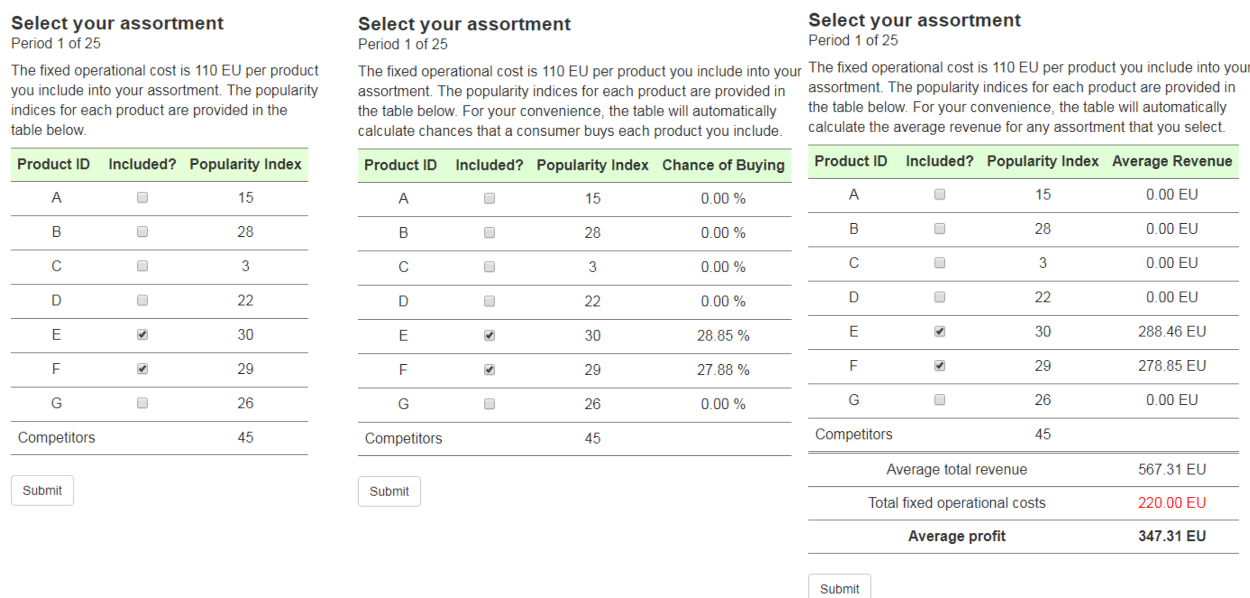

Select your assortment
Period 1 of 25

The fixed operational cost is 110 EU per product you include into your assortment. The popularity indices for each product are provided in the table below.

Product ID	Included?	Popularity Index
A	☐	15
B	☐	28
C	☐	3
D	☐	22
E	☑	30
F	☑	29
G	☐	26
Competitors		45

Submit

(a) No decision support (NS)

Select your assortment
Period 1 of 25

The fixed operational cost is 110 EU per product you include into your assortment. The popularity indices for each product are provided in the table below. For your convenience, the table will automatically calculate chances that a consumer buys each product you include.

Product ID	Included?	Popularity Index	Chance of Buying
A	☐	15	0.00 %
B	☐	28	0.00 %
C	☐	3	0.00 %
D	☐	22	0.00 %
E	☑	30	28.85 %
F	☑	29	27.88 %
G	☐	26	0.00 %
Competitors		45	

Submit

(b) Probabilities of buying (PB)

Select your assortment
Period 1 of 25

The fixed operational cost is 110 EU per product you include into your assortment. The popularity indices for each product are provided in the table below. For your convenience, the table will automatically calculate the average revenue for any assortment that you select.

Product ID	Included?	Popularity Index	Average Revenue
A	☐	15	0.00 EU
B	☐	28	0.00 EU
C	☐	3	0.00 EU
D	☐	22	0.00 EU
E	☑	30	288.46 EU
F	☑	29	278.85 EU
G	☐	26	0.00 EU
Competitors		45	
	Average total revenue		567.31 EU
	Total fixed operational costs		220.00 EU
	Average profit		**347.31 EU**

Submit

(c) Expected Profits (EP)

Figure EC.15 Decision stage screens for different levels of decision support information

Note: parameters shown are for Experiment 1 low optimal variety (LOV) condition

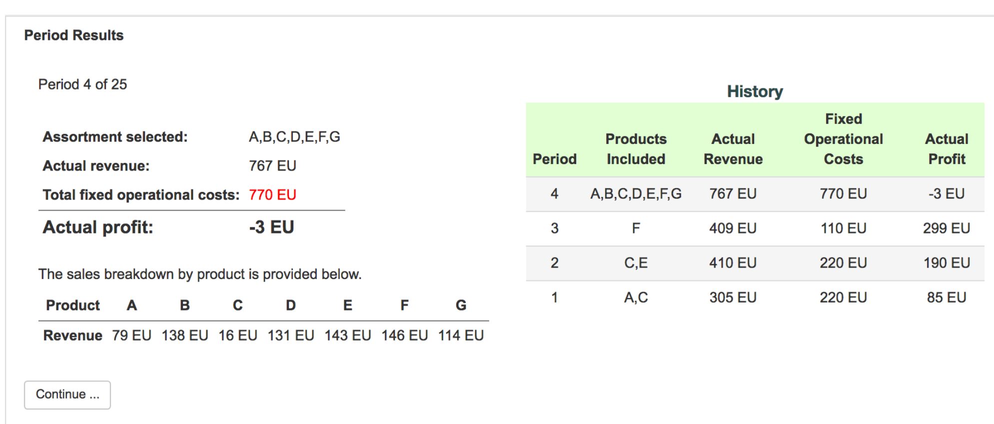

Period Results

Period 4 of 25

Assortment selected:	A,B,C,D,E,F,G
Actual revenue:	767 EU
Total fixed operational costs:	770 EU
Actual profit:	**-3 EU**

The sales breakdown by product is provided below.

Product	A	B	C	D	E	F	G
Revenue	79 EU	138 EU	16 EU	131 EU	143 EU	146 EU	114 EU

Continue ...

History

Period	Products Included	Actual Revenue	Fixed Operational Costs	Actual Profit
4	A,B,C,D,E,F,G	767 EU	770 EU	-3 EU
3	F	409 EU	110 EU	299 EU
2	C,E	410 EU	220 EU	190 EU
1	A,C	305 EU	220 EU	85 EU

Figure EC.16 Results stage screen

EC.7. Subject Instructions

Assortment Planning Game

In this experiment, you will assume a role of an assortment planner: that is, you will be picking the products that your company is offering on the market. You will be making this decision for 25 consecutive periods. The periods are *independent*, which means that your decisions and outcomes in one period will not affect outcomes in any other period. At the end of the game the profits in EU (Experimental Units) you have accumulated over all the periods will be converted to a US dollars amount and paid to your murk account as a HIT bonus.

How your sales are calculated

In each period there are 1,000 potential consumers. Each consumer buys exactly one unit: either from your assortment or from your competitors'.

Your marketing manager has conducted a survey and calculated a *popularity index* for each product that you can potentially include into your assortment. In addition to that, she has computed an aggregated popularity index for the products of your competitors. The chance that a consumer will buy a particular product is proportional to this product's popularity index.

Consider the example below: you can include two products into your assortment: product A and product B.

Product	Popularity Index
A	4
B	2
Competitors	6

If you include only product A in your assortment, the chance that a consumer buys it is:

- $4/(4+6) = 40\%$ for product A
- $6/(4+6) = 60\%$ for your competitors' products.

If you include both products A and B in your assortment, the chance that a consumer buys it is:

- $4/(4+2+6) = 33\%$ for product A
- $2/(4+2+6) = 17\%$ for product B
- $6/(4+2+6) = 50\%$ for your competitors' products.

In each period, each one of the 1,000 potential customers randomly chooses between the available products – yours and your competitors' assortments. The resulting sales for each product in your assortment have a normal[1] (bell-shaped) distribution with a mean of:

$$Average\ sales\ =\ Chance\ of\ buying\ \times\ 1{,}000$$

Your actual sales can be more or less than the average sales and will vary even if you select the same assortment in every period.

For example, if the chance that a customer buys product A is 40%, the average sales will be $40\% \times 1{,}000\ =\ 400$. The distribution of *actual sales* for this case is shown on the picture below.

Figure 1: Example of actual sales distribution for a single product

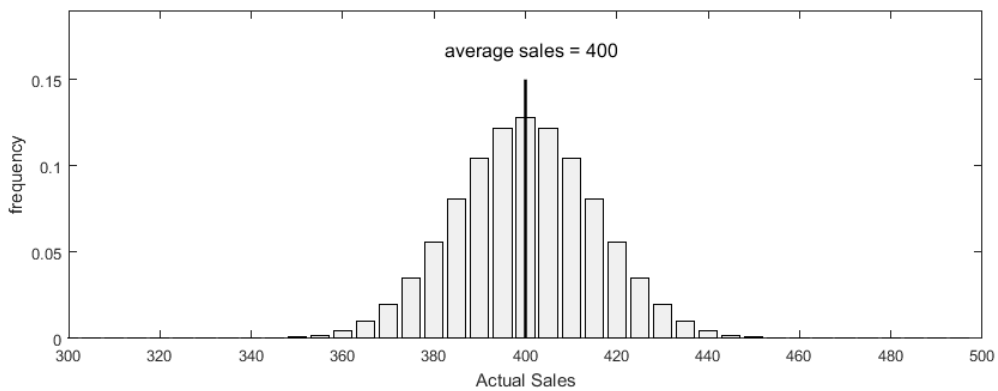

How you make money

You have an abundant supply of all products available immediately upon request, so you can satisfy any demand and you do not incur any inventory costs. However, there is a *fixed operational cost* for every product you include into your assortment. For example, if your assortment includes three products, and the operational cost is 50 EU, your total fixed operational costs are $50\ EU \times 3\ =\ 150\ EU$, regardless of your actual sales.

For each unit of product you sell, you earn **1 EU**. Your objective is to maximize your profit in each period, which is given by:

$$Profit\ = (1\ EU\ \times\ Actual\ Sales) - (Fixed\ Operational\ Cost\ \times\ Number\ of\ Products\ in\ the\ Assortment)$$

How your bonus payment is calculated

The game will end after 25 periods. You will then see your total profit (in EU), which is simply the sum of the profits you have accumulated across all 25 periods. Your total profit will be converted to a US dollars amount at the rate of **1000 EU = $1**.

[1] Actually, the distribution is binomial with a mean of $N \times p$ and a standard deviation of $\sqrt{N \times p \times (1-p)}$, where N is the number of potential consumers and p is the probability of buying. However, since N is very large, it very much resembles a normal distribution.