

From Customer Profiles to Dynamic Monitoring: Purchase-Sequence Signals and Fresh-Item Repurchase in Online Grocery Shopping

Jiaxi Wang

*College of Liberal Arts and Sciences, University of Illinois at Urbana-Champaign, Urbana, USA
jiaxiwang02@gmail.com*

Abstract. Customer summaries record how much a household has bought, not how its orders unfolded. We reconstruct ten-order histories for 1,477,553 target orders from 83,755 Instacart users and derive two sequence measures from each history: the coefficient of variation of interpurchase gaps, and the mean Jaccard overlap of consecutive fresh-item baskets. The outcome is the share of fresh items in a target basket that the household had bought before. Fractional-logit models split each measure into a customer mean and a window-level deviation. Timing irregularity is negative in both comparisons. Basket overlap is not: a one-standard-deviation increase in a customer's average overlap corresponds to 1.77 percentage points more repurchase, whereas the same increase relative to that customer's own average corresponds to 0.84 points less. The sign reversal survives eight-, ten- and twelve-order windows, strictly non-overlapping windows, a fixed complete-user sample, two opposite treatments of censored gaps, a narrower fresh-product definition, and an outcome whose reference period is aligned to the same ten orders. It also survives a customer fixed-effects linear model and a control for the preceding order's repurchase intensity. A customer's usual sequence and a recent departure from it are therefore distinct signals, and storing them as one similarity score discards the second.

Keywords: online grocery retail, interpurchase gaps, fresh-product repeat purchase, basket overlap, customer analytics

1. Introduction

Customer totals compress an order history and, in doing so, discard its sequence. Instacart records retain two kinds of information that such totals cannot recover: order dates reveal how purchases are spaced, and product identifiers show how the contents of one basket relate to the next. Both are consequential in an online setting, where the search and choice environment differs from store-based retail [1] and successive grocery baskets often overlap substantially [2].

Standard customer summaries typically begin with recency, frequency, and average purchase rate [3]. Research on transaction timing indicates that the spacing of purchases helps distinguish these cases [4-6]. Historical reorder rate captures repetition over an extended record, but it reveals little about whether adjacent baskets contain similar product sets.

The central issue is the level of analysis. A relationship observed across customers need not describe what happens when an individual customer's behavior changes. For the cross-customer comparison, each measure is represented by the customer's mean over the available ten-order windows. The within-customer comparison uses the difference between a particular window and that mean. Both terms enter the same fractional-logit specification.

To address these gaps, this study examines two sequence characteristics of online fresh-food consumption that have not been fully investigated: variation in purchase intervals and the degree of overlap of fresh-product sets between adjacent baskets.

2. Literature review

2.1. Purchase timing and basket-sequence structure

Purchase frequency and temporal regularity describe related but non-equivalent features of behavior. Gladly, Lemmens, and Croux link transaction spacing to both expenditure and dropout [6]. Platzer and Reutterer find that regularity improves inference about customer activity [4], and Reutterer, Platzer, and Schröder subsequently incorporate it into a parsimonious prediction framework [5].

Multicategory models make this point explicit by treating the basket as a joint purchase-incidence outcome [7]. Evidence from Instacart likewise shows that consecutive baskets belonging to the same household overlap much more than unrelated baskets [2]. The measure used in this study concerns that order-to-order continuity.

2.2. Persistent customer differences and within-customer variation

Repeated purchasing contains both persistent differences and shorter-lived departures from them. Related distinctions appear in research on state dependence, customer-relationship dynamics, and time-varying attrition, where stable heterogeneity and changing conditions are modeled separately [8-10]. Household panel data are particularly suitable for this purpose because the two sources of variation are observed within the same longitudinal record [11].

It is not established, however, that sequence variables have the same association at both levels. In the analysis, the between-customer component is represented by the mean of each customer's rolling-window values. The corresponding within-customer component is the deviation of a particular window from that mean.

3. Data and methods

3.1. Data and sample construction

We use the Instacart Online Grocery Shopping Dataset released in 2017. It contains anonymized records for more than 200,000 users and over three million orders. The order file establishes the sequence of purchases and gives the reported interval since the previous order. Product and department files identify basket contents. Their reordered flag indicates whether a product had already appeared in the user's purchase record.

Orders 1-10 form the first history window and order 11 is its target; the window then shifts forward by one order. Repeating this procedure across users produces 1,638,963 potential targets for 93,881 users. Targets lacking fresh items or any required model input are removed first, after which users represented by only one eligible target are excluded. The resulting panel contains 1,477,553 target orders contributed by 83,755 users.

In the baseline definition, fresh products come from four Instacart departments: produce, dairy eggs, meat seafood, and deli. To assess whether the results depend on this classification, a sensitivity specification retains only produce and meat seafood.

3.2. Outcome, sequence measures, and controls

For customer i and target order t , let N_{it}^F denote the number of fresh products purchased and R_{it}^F the number marked reordered = 1. When $N_{it}^F > 0$, Equation (1) defines the dependent variable. Because the reordered indicator is determined using purchases prior to t , the numerator includes only fresh products already found in the customer's pre-target history. A product that appears only after the target order cannot affect Y_{it} .

$$Y_{it} = \frac{R_{it}^F}{N_{it}^F}, \quad N_{it}^F > 0 \quad (1)$$

The restriction $N_{it}^F > 0$ is part of the estimand, not merely a sample-cleaning decision. If a target order contains no fresh product, a zero would mix category nonparticipation with a basket containing fresh products but no familiar items. Y_{it} instead records the previously purchased share once the customer has entered the fresh category in that order. Purchase-incidence models make the same conceptual separation between whether a category is chosen and what enters the basket after choice [7].

Purchase timing is summarized by the coefficient of variation across nine adjacent gaps in a ten-order history. With mean gap $\mu_{g,it}$ and standard deviation s_{it} , Equation (2) defines the timing measure. This ratio separates gap dispersion from average purchasing pace.

$$T_{it} = \frac{s_{it}}{\mu_{g,it}} \quad (2)$$

Instacart records all intervals of 30 days or more as 30; their true lengths are therefore unknown. Such values make up 3.37% of otherwise usable gaps. The baseline timing measure excludes them, while a sensitivity model retains the recorded value.

Basket continuity is based on consecutive order pairs. For pair m before target t , j_{itm} is the Jaccard overlap of the two fresh-item sets. It equals zero when both orders contain fresh items but no item overlaps. Pairs with no fresh item in either order are omitted because two empty sets contain no composition information. Equation (3) averages the remaining M_{it} pair values.

$$J_{it} = \frac{1}{M_{it}} \sum_{m=1}^{M_{it}} j_{itm} \quad (3)$$

Five pre-target characteristics enter as controls. They are average basket size, mean interpurchase interval, fresh-product share, historical fresh reorder rate, and log cumulative prior orders. We standardize variables within the sample used by each specification. Consequently, differences in raw coefficients across alternative samples may partly reflect changes in the sample standard deviation and are interpreted descriptively.

3.3. Empirical strategy and model specification

Because Y_{it} is fractional and includes observations at $Y_{it} = 0$ and $Y_{it} = 1$, we estimate fractional-logit models [12]. Equation (4) presents the pooled version. H_{it} denotes the conditional purchase

history, T_{it} and J_{it} are the two sequence measures, X_{it} contains the controls, and $\Lambda(\cdot)$ is the logistic cumulative distribution function.

$$E[Y_{it} | H_{it}] = \Lambda(\alpha + \beta_1 T_{it} + \beta_2 J_{it} + \gamma^\top X_{it}) \quad (4)$$

Equation (4) is reported as a pooled benchmark. It does not distinguish a difference between two customers from a change across windows belonging to one customer. Since a user contributes several target orders, standard errors are clustered at the customer level. Adjacent target orders also share nine of the ten orders in their histories, meaning that the regressors are themselves autocorrelated by construction; clustering addresses dependence in the errors rather than in the design, and the effective sample is accordingly smaller than the nominal count.

The panel allows the two comparisons to be separated. Suppose customer i contributes T_i usable targets and S_{it} is any time-varying regressor. Its observed customer mean is $\mu_{S,i}$. We define the within-customer deviation as in Equation (5), so the second term records the position of window t relative to that customer's own average.

$$\mu_{S,i} = \frac{1}{T_i} \sum_{t=1}^{T_i} S_{it}, \quad S_{it}^W = S_{it} - \mu_{S,i} \quad (5)$$

In this parameterization, $\mu_{S,i}$ carries the between-customer comparison and S_{it}^W the within-customer comparison. Including unit means implements a correlated-random-effects approach to persistent heterogeneity [13, 14]. We apply the split to every time-varying control and standardize the two sets of components separately. Means are based on the targets observed for each user, which permits an unbalanced-panel correlated-random-effects specification [15].

Substituting customer means and window deviations for the time-varying regressors yields Equation (6), which is the specification used for the main estimates.

$$E[Y_{it} | H_{it}] = \Lambda\left(\alpha + \beta_1 \mu_{T,i} + \beta_2 T_{it}^W + \beta_3 \mu_{J,i} + \beta_4 J_{it}^W + \delta_1^\top \mu_{X,i} + \delta_2^\top X_{it}^W\right) \quad (6)$$

The parameter interpretation in Equation (6) follows from this decomposition. β_1 and β_3 compare customers on their mean timing irregularity and mean basket similarity; β_2 and β_4 instead measure departures within a customer. Inference is based on customer-clustered standard errors. To express the results on the outcome scale, average partial effects are computed by evaluating the derivative of the conditional mean for every observation and then averaging those derivatives. The clustered covariance matrix enters the Delta-method calculation of the APE standard errors and 95% confidence intervals.

3.4. Robustness analyses

Robustness checks vary the functional form, window design, sample, and measurement. We estimate a customer fixed-effects linear model, replace the ten-order history with eight- and twelve-order windows, impose non-overlapping windows, and use a fixed complete-user sample. For timing, the recorded 30-day value is alternately retained or the affected histories are removed. Fresh products are also narrowed to produce and meat seafood. We additionally redefine the outcome so that a fresh item is counted as repurchased only if it appears among the preceding ten orders, aligning the reference period of the outcome with that of the sequence measures. The final diagnostic uses adjacent eligible targets, reproduces the baseline on that common sample, and then controls for the

previous target's conditional fresh-repurchase intensity. Table 1 reports descriptive statistics for the estimation panel on the original scale.

Table 1. Descriptive statistics for the 10-order rolling panel

Variable	Mean	SD
Conditional fresh-repurchase intensity	0.809	0.261
Purchase timing irregularity	0.591	0.230
Successive fresh-basket similarity	0.194	0.150
Average basket size	10.419	5.716
Mean interpurchase interval	7.234	3.495
Fresh-product share	0.543	0.177
Historical fresh reorder rate	0.752	0.179
Cumulative prior orders	28.306	18.044

Notes: Original-scale statistics for 1,477,553 targets from 83,755 users. Cumulative prior orders enters the regression in logarithms.

4. Results

4.1. Outcome and measure structure

Among 1,638,963 potential targets, 152,592 contain no fresh product, 60,095 contain fresh products but no previously purchased item, and 1,426,276 contain at least one reordered fresh item. Coding orders outside the fresh category as zero would create 71.74% of all zero cases. The estimand therefore concerns repurchase intensity conditional on fresh-category participation.

The measures also vary within users. Overall and within-customer standard deviations are 0.150 and 0.070 for basket similarity, and 0.230 and 0.159 for timing irregularity. Among repeated users, similarity changes for 96.22% and timing irregularity for 99.33%. Historical fresh reorder rate correlates more strongly with similarity between customers ($r = 0.577$) than within them ($r = 0.257$). The maximum variance inflation factor is 3.58 (Table 2).

Table 2. Variance inflation factors by level of variation

Variable	Between VIF	Within VIF
Basket similarity	1.90	1.22
Timing irregularity	1.35	1.03
Average basket size	1.15	1.09
Mean interpurchase interval	2.83	1.03
Fresh share	1.14	1.06
Historical fresh reorder rate	2.78	1.69
Log cumulative prior orders	3.58	1.54

Notes: Computed over all between- and within-customer regressors in the main sample, the maximum is 3.58.

4.2. Main within-between estimates

Table 3 presents the pooled benchmark. With purchase-history controls, a one-standard-deviation increase corresponds to average partial effects of -0.86 percentage points for timing irregularity and +2.44 for basket similarity. Table 4 separates the customer means from the within-customer deviations that the pooled estimates combine.

Table 3. Pooled fractional-logit estimates

Standardized variable	Coefficient	Clustered SE	APE (pp)
Timing irregularity	-0.0596	0.0017	-0.86
Basket similarity	0.1685	0.0026	2.44
Average basket size	0.0747	0.0017	1.08
Mean interpurchase interval	-0.1084	0.0018	-1.57
Fresh share	0.0876	0.0018	1.27
Historical fresh reorder rate	0.4143	0.0025	5.99
Log cumulative prior orders	0.0434	0.0023	0.63

Notes: Estimated on 1,477,553 orders from 83,755 users. Regressors are standardized, standard errors clustered by user, and average partial effects are percentage-point changes for a one-standard-deviation increase. Intercept omitted.

Table 4. Within-between fractional-logit estimates of conditional fresh-repurchase intensity

Panel A. Fractional-Logit Coefficients		
Variable	Between-customer	Within-customer
Timing irregularity	-0.0290 (0.0018)	-0.0152 (0.0016)
Basket similarity	0.1248 (0.0027)	-0.0593 (0.0022)
Average basket size	-0.0084 (0.0016)	0.0667 (0.0018)
Mean interpurchase interval	-0.0440 (0.0026)	-0.0483 (0.0016)
Fresh-product share	0.0008 (0.0017)	0.1047 (0.0019)
Historical fresh reorder rate	0.7327 (0.0032)	0.0255 (0.0025)
Log cumulative purchase history	-0.0372 (0.0025)	0.1439 (0.0027)
Panel B. Average Partial Effects (percentage points)		
Variable	Between-customer [95% CI]	Within-customer [95% CI]
Timing irregularity	-0.41 [-0.46, -0.36]	-0.22 [-0.26, -0.17]
Basket similarity	1.77 [1.69, 1.84]	-0.84 [-0.90, -0.78]

Notes: Estimated on 1,477,553 orders from 83,755 users. Customer means and within-customer deviations enter the same regression and are standardized separately. Customer-clustered standard errors in parentheses. Average partial effects are percentage-point changes for a one-standard-deviation increase, with 95% confidence intervals in brackets. Intercept omitted.

Timing irregularity has a negative coefficient in both comparisons. For the customer-mean component, $\beta = -0.0290$ and the average partial effect is -0.41 percentage points. For a deviation

from the customer mean, $\beta = -0.0152$ and the effect is -0.22 points.

Across customers, the coefficient on mean basket similarity is 0.1248 and its average partial effect is +1.77 percentage points. Within a customer, the coefficient is -0.0593, corresponding to -0.84 points. In the latter comparison, repurchase in the next target is slightly lower. Historical fresh reorder rate remains controlled for in the equation, together with the other controls, meaning that this contrast concerns continuity between neighboring baskets over and above the customer's accumulated tendency to reorder. One within-customer standard deviation equals 0.070 Jaccard points. Deviations at the 10th and 90th percentiles are -0.075 and 0.079, respectively, placing the -0.84-percentage-point effect within a range of changes that occurs regularly in the panel. Relative to a mean outcome of 0.809, the association is small in absolute terms. A single aggregate similarity score would not preserve the distinction in sign.

4.3. Robustness and sensitivity checks

Table 5 groups the checks by what is allowed to vary. Panel A changes the functional form on the baseline sample; Panels B through D vary the window, the sample, the measurement, and the treatment of the preceding outcome. Panel A and the baseline share the bounded outcome and the standardized within regressors, so the two are directly comparable on the outcome scale. Coefficients estimated on other samples are interpreted descriptively, because standardization is repeated within each sample.

Table 5. Robustness and measurement sensitivity of the within-between pattern

Specification	N (users)	Timing: Btw	Timing: Wtn	Basket: Btw	Basket: Wtn
Panel A. Functional-form check					
Customer fixed-effects linear	1,477,553 (83,755)	—	-0.0018 (0.0002)	—	-0.0084 (0.0003)
Panel B. Window design and sample					
8-order window	1,651,489 (91,245)	-0.0296 (0.0017)	-0.0155 (0.0015)	0.1253 (0.0025)	-0.0553 (0.0019)
10-order window (baseline)	1,477,553 (83,755)	-0.0290 (0.0018)	-0.0152 (0.0016)	0.1248 (0.0027)	-0.0593 (0.0022)
12-order window	1,317,547 (72,564)	-0.0279 (0.0019)	-0.0139 (0.0018)	0.1243 (0.0029)	-0.0600 (0.0024)
Fixed complete-user subset	942,628 (53,376)	-0.0295 (0.0022)	-0.0131 (0.0020)	0.1254 (0.0034)	-0.0583 (0.0027)
Strict non-overlap	147,316 (43,445)	-0.0212 (0.0050)	-0.0168 (0.0045)	0.0863 (0.0077)	-0.0359 (0.0058)
Panel C. Measurement sensitivity					

Table 5. (continued)

30-day coded gaps included	1,477,553 (83,755)	-0.0291 (0.0017)	-0.0174 (0.0017)	0.1237 (0.0027)	-0.0593 (0.0022)
No 30-day boundary gaps	1,187,816 (60,451)	-0.0242 (0.0020)	-0.0170 (0.0018)	0.1311 (0.0031)	-0.0539 (0.0024)
Narrow fresh definition	812,127 (50,175)	-0.0290 (0.0024)	-0.0153 (0.0024)	0.1549 (0.0035)	-0.0537 (0.0031)
10-order-aligned outcome	1,477,553 (83,755)	-0.0382 (0.0019)	-0.0123 (0.0013)	0.2562 (0.0028)	-0.0731 (0.0019)
Panel D. Lagged-outcome diagnostic					
Same-sample baseline	1,309,584 (75,701)	-0.0264 (0.0019)	-0.0154 (0.0017)	0.1499 (0.0029)	-0.0692 (0.0023)
Lagged outcome controlled	1,309,584 (75,701)	-0.0122 (0.0014)	-0.0155 (0.0017)	0.1323 (0.0023)	-0.0689 (0.0023)

Notes: Panel A is a customer fixed-effects linear model, Panels B-D are within-between fractional logits with user-clustered standard errors in parentheses. Panel A shares the baseline sample and standardized within regressors. Alternative samples are standardized separately, so raw coefficients across samples are descriptive and the direction rather than the magnitude is the comparison that carries. The 10-order-aligned outcome row is the exception: it holds the baseline sample and scaling fixed and changes only the reference period of the outcome, and is therefore directly comparable with the baseline row. Panel D holds the adjacent-target sample fixed before adding lagged-outcome components.

The result is not driven by the functional form. A customer fixed-effects linear model returns -0.18 percentage points for timing and -0.84 for basket similarity, against fractional-logit effects of -0.22 and -0.84 on the same sample. The basket estimates agree to two decimal places despite different link functions and different treatments of persistent heterogeneity.

Nor is the window length. Eight-, ten- and twelve-order histories preserve both signs, and the basket within coefficient moves only between -0.0553 and -0.0600 across the three. Restricting the panel to users present throughout leaves the estimates close to the baseline.

The window overlap requires separate discussion. In the baseline, adjacent targets share nine of their ten history orders, so successive observations are far from independent. Imposing strict non-overlap cuts the panel to 147,316 observations from 43,445 customers and reduces the mean number of targets per customer from 17.64 to 3.39; the basket within coefficient attenuates from -0.0593 to -0.0359. Because this sample is restandardized, the direction rather than the magnitude is the comparison that carries; the two coefficients are not expressed on a common scale, and the attenuated value should not be read as a bound on the baseline.

Measurement choices do not generate the pattern either. Retaining the censored 30-day gaps, removing every history that contains one, and narrowing fresh products to produce and meat seafood all preserve both signs.

One measurement check is directly comparable to the baseline, because it holds the sample and the scaling fixed and changes only the outcome. Aligning the reference period of repurchase to the same ten orders used for the sequence measures leaves 1,477,553 observations from 83,755 customers, and both patterns persist. The basket coefficient is 0.2562 between customers and -0.0731 within them, against 0.1248 and -0.0593 in the baseline; timing remains negative at both

levels. Removing the mismatch between the two reference periods therefore does not weaken the contrast; on this sample it is somewhat larger.

The final diagnostic uses the 1,309,584 adjacent targets contributed by 75,701 users. On that common sample the basket within coefficient is -0.0692; adding the previous target's conditional fresh-repurchase intensity moves it to -0.0689, a change of 0.5%. The immediately preceding outcome therefore does not absorb the negative within-customer association.

The same control affects the two between-customer terms very differently. Timing irregularity falls from -0.0264 to -0.0122, while basket similarity moves only from 0.1499 to 0.1323. More than half of the cross-customer timing association is thus shared with the household's accumulated repurchase record, whereas the cross-customer basket association is largely distinct from it. This asymmetry was unexpected, and it strengthens rather than qualifies the paper's claim: the two sequence measures differ not only in the sign of their within-customer components but in how much of their between-customer variation is already contained in conventional history summaries. The exercise remains a diagnostic and is not interpreted as a dynamic causal model.

5. Discussion

Two patterns emerge from Table 4. Timing irregularity is negative in both columns, which extends work on purchase regularity beyond average frequency [4-6]. Basket similarity is not: it is positive across customers and negative within a single customer's history. The association is therefore level-specific, and the distinction between persistent customer differences and recent departures from them [8-11] cannot be carried by one pooled coefficient.

The magnitude deserves a plain statement. Against a mean outcome of 0.809, a within-customer effect of -0.84 percentage points is a relative change of roughly one percent. One within-customer standard deviation equals 0.070 Jaccard points, and the tenth and ninetieth percentiles of the deviation are -0.075 and 0.079, so the effect corresponds to a movement that occurs routinely in the panel rather than to an extreme case. The claim advanced here is accordingly about sign and about feature design, not about the size of an intervention: a pooled coefficient of +2.44 percentage points would attribute an increase in repurchase to a customer whose baskets are in fact drifting away from their own pattern.

More than one consumer process could produce the basket result, and the design does not separate them. Persistently high similarity is consistent with a habitual repertoire sustained by stable purchase cues. A short run of unusually repetitive baskets could instead precede a return to variety, or track the replenishment cycle of household inventories. These accounts are not mutually exclusive, and none of them is tested here.

The decomposition translates directly into a feature specification. For each sequence measure, the customer mean can be stored as a profile field and the most recent window deviation as a monitoring field. Timing irregularity may be read the same way in both, since its components share a sign. Basket similarity must be kept as two columns, because collapsing them removes the only component that changes over the short run. Both inputs are available from routine order logs, and neither requires price, promotion, or demographic data. In operation, the means can anchor segmentation and benchmark groups, while the deviations can place customers into review queues when their ordering rhythm or basket continuity shifts. Category entry should remain separate from repurchase intensity after entry, so that absence from the fresh category is not recorded as weak repeat purchasing.

Two limitations remain, in descending order of their relevance to the central claim. A third, concerning measurement alignment, was tested rather than left open.

The first concerns the within-customer basket estimate. If the similarity measure carries window-level noise, a window that happens to be measured high will tend to be followed by a lower one whatever the household does, and mean reversion would produce a negative coefficient mechanically. The lagged-outcome diagnostic in Section 4.3 does not settle this, because it addresses persistence in the outcome rather than measurement error in the regressor. A direct test is available but is not attempted here: it would involve computing basket similarity separately on the first and last five orders of each window and comparing the two within-customer coefficients. If the association were mechanical, the late-window version should be substantially the larger of the two.

The second concerns the operational recommendation. The paper argues from the sign reversal that two columns should be stored rather than one, but it does not demonstrate that two columns predict better. A held-out comparison, splitting customers rather than orders, would test that directly and is left to further work.

Measurement alignment was the third concern and is now addressed in Table 5. The reordered flag is defined against a customer's entire prior record, whereas the sequence measures are computed on a ten-order window, so for target orders late in a history the two reference periods do not coincide. Recomputing the outcome on the same ten orders removes the mismatch without changing the sample or the scaling, with both patterns persisting. What remains is that the aligned outcome is a narrower construct: it records repurchase of items seen recently rather than repurchase of items ever seen, and the two are not the same quantity.

Finally, all estimates are associative. They may inform segmentation, monitoring, and the selection of customers for closer review, but they do not identify the response to a retailer action. Contemporaneous measures of price, promotion, quantity, and stock availability would be required for that question.

6. Conclusion

Purchase sequences retained in transaction logs change how conventional customer-history variables should be interpreted. In the rolling Instacart panel, timing irregularity is negatively associated with conditional fresh-item repurchase both across customers and within the same customer. Basket similarity has a level-specific association: the customer-mean component is positive, whereas the deviation from that mean is negative.

A customer mean summarizes the usual purchase sequence and is suited to profiling. The most recent window deviation records movement away from that sequence and is suited to monitoring.

These applications are associative rather than causal. The measures may inform segmentation, monitoring, and the selection of customers for closer review, but they do not estimate the response to a retailer intervention. Addressing that question would require contemporaneous measures of price, promotion, quantity, and stock availability.

References

- [1] Grewal, D., Roggeveen, A. L., & Nordfält, J. (2017). The future of retailing. *Journal of Retailing*, 93(1), 1-6.
- [2] Chintala, S. C., Liaukonytė, J., & Yang, N. (2024). Browsing the aisles or browsing the app? How online grocery shopping is changing what we buy. *Marketing Science*, 43(3), 506-522.
- [3] Fader, P. S., Hardie, B. G. S., & Lee, K. L. (2005). Counting your customers the easy way: An alternative to the Pareto/NBD model. *Marketing Science*, 24(2), 275-284.
- [4] Platzter, M., & Reutterer, T. (2016). Ticking away the moments: Timing regularity helps to better predict customer activity. *Marketing Science*, 35(5), 779-799.

- [5] Reutterer, T., Platzer, M., & Schröder, N. (2021). Leveraging purchase regularity for predicting customer behavior the easy way. *International Journal of Research in Marketing*, 38(1), 194-215.
- [6] Gladys, N., Lemmens, A., & Croux, C. (2015). Unveiling the relationship between the transaction timing, spending and dropout behavior of customers. *International Journal of Research in Marketing*, 32(1), 78–93.
- [7] Manchanda, P., Ansari, A., & Gupta, S. (1999). The shopping basket: A model for multicategory purchase incidence decisions. *Marketing Science*, 18(2), 95–114.
- [8] Dubé, J. P., Hitsch, G. J., & Rossi, P. E. (2010). State dependence and alternative explanations for consumer inertia. *RAND Journal of Economics*, 41(3), 417–445.
- [9] Netzer, O., Lattin, J. M., & Srinivasan, V. (2008). A hidden Markov model of customer relationship dynamics. *Marketing Science*, 27(2), 185–204.
- [10] Bachmann, P., Meierer, M., & Näf, J. (2021). The role of time-varying contextual factors in latent attrition models for customer base analysis. *Marketing Science*, 40(4), 783–809.
- [11] van Heerde, H. J., & Dekimpe, M. G. (2024). Household and retail panel data in retailing research: Time for a renaissance? *Journal of Retailing*, 100(1), 104–113.
- [12] Papke, J. L., & Wooldridge, J. M. (1996). Econometric methods for fractional response variables with an application to 401(k) plan participation rates. *Journal of Applied Econometrics*, 11(6), 619–632.
- [13] Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica*, 46(1), 69–85.
- [14] Bell, A., & Jones, K. (2015). Explaining fixed effects: Random effects modeling of time-series cross-sectional and panel data. *Political Science Research Methods*, 3(1), 133–153.
- [15] Wooldridge, J. M. (2019). Correlated random effects models with unbalanced panels. *Journal of Econometrics*, 211(1), 137–150.