A Researcher's Checklist for Working with Sales Data to Evaluate Healthy Retail Interventions

Released by:
Dr. Molly De Marco, Ph.D., MPH[1]
Christina Chauvenet, M.Sc., MSPH[1]
Leah Chapman, MPH[1]
Danton Noriega-Goodwin, M.Sc., MA[2]

[1]UNC Center for Health Promotion and Disease Prevention
[2]Duke University, Sanford School of Public Policy

June 5, 2017

Introduction

A growing number of researchers are interested in using retail sales data to evaluate the effectiveness of healthy retail interventions. Sales data can be used to examine whether in-store promotions targeting certain food and beverage products affects sales by comparing the sales data from before the intervention to during and after the intervention. A 2013 systematic review of 33 healthy retail interventions in grocery stores found that 13 studies reported using objective store sales data as an outcome metric.[1] Use of sales data has numerous benefits over other outcome metrics, the primary advantage being that sales data is an objective measure of customers' purchasing decisions strengthening study internal validity.[2] Other evaluation measures, such as customer intercept surveys or store manager interviews, can provide valuable insight, but as self-report measures, they are subject to bias.[2]

While there are clear advantages to using sales data, there are also challenges involved in procuring, organizing, storing, and analyzing such data. This brief aims to provide an overview of key considerations for researchers who wish to use sales data to evaluate the effectiveness of healthy retail interventions. This brief discusses key considerations for identifying the research question, forming a partnership with retailers, and data collection and analysis.

This brief looks explicitly at *sales* data, or data obtained by the retailer at the point of sale, as opposed to *purchase* data, which is data obtained by having individuals submit receipts or other proof of purchases directly to researchers. While we focus on using sales data to evaluate in-store interventions, sales data could also be used to learn about retail behavior without implementing an intervention, and the same recommendations would apply.

Steps to Take

1. Identifying the Research Question

Identifying the research question is integral to developing a partnership with retailers. The research question will largely depend on the research setting and availability of data in that research setting. Depending on the ultimate goal of the researchers, research questions may be developed based on the research setting and availability of data in an already existing partnership or retail setting, or the research question may drive the development of partnerships.

In the first scenario, the researcher may already have developed a retailer relationship or have a particular type of research setting in mind. For example, if a researcher is interested in food purchases in rural corner stores in a sparsely populated county, the researcher will have limited potential partners. In these settings, the research question must be developed with the availability of data in these settings in mind. The research question should be developed in conjunction with the partner retailer so that the research question matches

the availability of data. For example, many small retailers do not have access to sophisticated point-of-sale systems that allow for the capture of data on the sale of individual items.

The development of this research question can and should be an iterative process. If the data available is not sufficient to answer key research questions of interest, the researcher may consider investing in systems to help obtain the level of data desired. The researcher must identify what data is available or plausible to collect and does not create an undue burden on the retailer. The research question must map on to the availability of data to be successful. As another example, if a store does not have a loyalty card program, then customers cannot be linked to repeated sales and inferences about sales cannot be made to the individual level.

If researchers have a broader availability of research questions, then the research question may also drive the development of partnerships. For example, if the researcher is specifically interested in looking at differences in sales of fruits and vegetables by race/ethnicity, but is not constricted to a particular geographic area, the researcher should seek out a retail partner who has demographic information through loyalty card programs.

The development of successful partnerships between retailers and researchers involved an iterative development of a research question based on both researcher and retailer needs. As discussed below, minimizing burden on the retailer is essential to a successful partnership, as is mapping the research question on to the type of data available.

2. Building a Partnership

Sales data provides opportunities and challenges for researchers.[3] Retail partners may be reluctant to work with researchers due to concerns regarding the perceived onerousness of the time commitment for their personnel, the burden of data collection, uncertainty about whether their investment will pay off, and whether the presence of researchers and the implementation of new strategies could negatively impact their sales.[4] Other concerns include crowding, shoplifting, or promoted items that do not sell.[3] It is therefore vital for researchers to develop strong, trustful relationships with their retail partners before requesting sales data. Researchers should explain potential benefits of partnership to retailers, such as providing the researchers with an in-depth analysis of their sales data. Researchers should accommodate requests for more information from retailers. For more information on recruiting retail partners please review the BECR brief entitled: "*What's in it for Retailers? Establishing Partnerships with Food Retailers to Conduct Healthy Food Choice Research.*"[4] Another useful resource is the Healthy Eating Research-Nutrition and Obesity Policy Research and Evaluation Network (NOPREN) Healthy Food Retail Working Group's *Grocery Retailer Academic Collaborative (GRAC) Guidelines.*[5]

A. Formalizing the Partnership and Determining Availability of Data

Once a retailer has agreed to provide sales data, several steps remain. First, the retailer and researcher should both sign a "Memorandum of Understanding" (MOU), which is a nonbinding agreement that outlines the terms and details of the partnership, including each parties' requirements and responsibilities.[6] Additionally, the researcher and retailer should determine how and when sales data will be delivered and received. An ideal dataset will include:

- ☐ categorization of transactions per customer (i.e. panel data format) for loyalty card customers (if available) or at least transactions by number.
- ☐ the time and date of each transaction
- ☐ itemized descriptions of type and quantity of units purchased including unique code (UPC or PLU)
- ☐ price of units purchased
- ☐ total store sales
- ☐ customer demographic data
- ☐ store demographic data (square footage, number of checkout lanes, number of employees, store hours, for example)

- ☐ inventory of store items and updates as inventory changes to allow tracking of newly introduced items
- ☐ trends of sales data via sales data of the same detail as above data from the past year

It is always preferable to have data from individual sales as well as storewide sales when possible. Having individual transaction data allows the researcher to examine two levels of data simultaneously, and the researcher can explore individual transaction patterns within store transaction patterns. For example, with transaction-level data the researcher can answer research questions about what items are purchased together.

However, this type of data may not be available in all settings. If the retailer cannot provide this level of detail on sales due to technological challenges, the researcher could consider investing in equipment for the store (such as purchasing new registers or new software for existing registers) to help the store improve their sales tracking. The ability to cover such costs will depend on a researcher's funding and is something that should be confirmed prior to suggesting this to the retail partner. The prioritization of improving the quality of data available will help improve the quality of the study and will bring value to the partnership for the retailer.

B. Securely Storing the Data

Retailers are often reticent to provide sales data because they are concerned this data could be used by their competitors to gain an advantage, an understandable concern given the tight margins of food retailers. Because of this very salient concern to retailers, it is essential to explain how the research team will keep a retailer's sales data and other proprietary information secure and confidential.

The following assurances can be offered when proposing a research partnership with a retailer:

1. Offer the option of a Data Use Agreement (DUA) or Memorandum of Understanding (MOU) that spells out the scope of the project, security measures, data confidentiality, and rights of all parties. An MOU would be broader and detail the resources brought to bear by both partners plus a timeline.

2. Assure data security and provide details on where the data would be maintained.

3. Offer the retailer anonymity in any publications or interviews regarding the research with the option to identify them if doing so provides benefit to the retailer. The decision to publish identifiable information should be the decision of the retailer.

4. Offer the retailer the opportunity to review products of the research, not for approval, but as an effort to ensure that the researcher complies with the DUA and any other preapproved requests and also as a matter of courtesy and good will.

3. Data Collection and Data Analysis Considerations

One advantage of working with retailers to obtain sales data is that most retailers already collect sales data. Retailers will already have systems in place to track store purchases and categorize them. Before suggesting strategies for tracking of sales data, researchers should familiarize themselves with the current tracking system for in-store purchases. To avoid creating a burden for the retailer, the researcher should aim to add desired tracking elements on to their current system rather than creating new systems when possible. There is variability in the level of detail to which stores track their purchases. Below, we use the example of fruits and vegetable purchases to demonstrate analytic considerations, though these considerations can be generalized to other foods and beverages.

A. Analyzing Sales Data: Absolute Values or Relative Changes to Sales

Whether a researcher analyzes absolute values of the sales of a particular item or looks at relative changes in sales in relation to the sales of other items is dependent on the research question. For example, if the research question is "Does the promotion of fruits and vegetables increase the sales of produce?", information about

sales of other items may not be needed. However, if the research question is "Does the promotion of fruits and vegetables improve the overall healthfulness of items purchased?" then information on other items purchased during each transaction are needed.

Here, it is also important to ascertain what the expected change or outcome of an intervention is and tailor analysis accordingly. For example, if the intervention is solely a fruit and vegetable price promotion, then the proportion of produce as a part of the total sale is likely the best data for analysis. If the intervention includes nutrition education pamphlets about the value of swapping unhealthier snacks (e.g. candy bars) for fruits and vegetables, then an analysis of the proportion of the unhealthy to healthy food items purchased is merited.

B. Categorization of Data

Determining how sales data is collected at point of sale is important to determine at the onset of the intervention. Retailers may or may not already aggregate sales by specific type of item (e.g. sugar sweetened beverages, fruits and vegetables, meats, prepared food from in-store counters, etc.). In the fruit and vegetable example, researchers would want to ascertain whether sales of fruit and vegetables are already aggregated together. Particularly in small retail settings, sales data may only track fruit as "loose fruit", for example, if fruit prices are the same for multiple items. Likewise, all sodas may be tracked by size rather than type of soda (e.g. 20-ounce soda may refer to both Pepsi and Diet Pepsi).

Researchers should work with retailers to get the most disaggregated level of data possible. Researchers can then manipulate the data to create categories to explore food group types. Even if the original research question is about an aggregate food group (e.g. fruits and vegetable consumption), it is still best practice to obtain disaggregated data (by item like bananas instead of by category like fruit) in case further lines of research using this data would like to look at variation within a category. Data collected at a disaggregated level can easily be aggregated into food groups, but data that is already aggregated cannot be transformed back into individual items.

Depending on the research question, linking nutritional information to items in store may be helpful for categorization of items into item category and/or analysis of the nutrient profile of sales.

C. Linking Nutritional Information and Creating Food Group Categories

If the researcher wants to conduct analysis of nutritional information, matching sales data with food composition tables will likely be necessary through use of the United States Department of Agriculture's (USDA) National Nutrient Database for Standard Reference (USDA-SR).[7] This comprehensive database provides nutrient profiles for over 184,000 foods and branded food products. The database is free and available to the public.

The researcher should also consider what forms of the product they would want to track; for example, is the researcher only interested in fresh fruits and vegetables sales, or should frozen and canned fruit and vegetable sales be included in these sales? What about fruit or vegetable juices?

Analytic considerations for food composition analysis are out of the scope of this brief, but researchers should consult Brinkerhoff et al and Pennington et al (see references) for additional guidance on nutritional mapping to USDA SRs and creating food groups.[8, 9]

D. Monitoring Sales and the In-Store Environment

It is essential for the researcher to monitor sales data and the retail environment before and after the intervention is implemented. The ability to retrospectively examine sales patterns and stocking changes will again depend on the system of tracking sales the retailer already has in place. If sales records are available, we recommend obtaining the last 12 months of sales data to be able to see fluctuation in sales over time, on different days of the week, during different seasons, and during holidays.

Whether or not sales data from the last year is available, we recommend monitoring sales and the retail environment for at least 60 days before and after the intervention. Monitoring sales data post-intervention will allow for the examination of whether any behavior change is sustained after the intervention is removed. Changes to the in-store environment might include the rearrangement of store items on a particular shelf, promotions of other products besides your targeted item, and changes to stocking.

It is important to track stocking changes throughout the intervention as new items will affect sales patterns. Researchers should plan to have research assistants monitor the sales environment on a daily basis, and should ask retailers to inform the research team when item stocking changes or when there are changes to how the store is organized. Because retail environments change frequently, research assistants are particularly useful in monitoring the in-store environment to lessen the burden on the retailers. Research assistants should monitor what other promotions are going on in the store, and should be mindful that these promotions will change, likely on a daily basis. The changes to the in-store environments should be included as control variables in regression analyses of the outcome of interest.

E. Changes to Factors External to the Retail Outlet

As with the in-store environment, it is also critical to monitor the factors external to the retailer that may affect sales of foods and beverages. The built environment around the retail outlet can influence in-store sales. One built environment factor that may affect store purchases is other retail outlets surrounding the store. For example, in the fruit and vegetable example, if a farmers' market opened across the street from the same grocery store where the intervention was taking place, there may be a decrease in fruit and vegetable sales as a proportion of total sale, though this change may not reflect an ineffective intervention but rather a shift of sales from one retail outlet to another.

There may also be changes to accessibility to the store due to factors in the built environment. For example, road construction around the store may reduce store sales, whereas the addition of sidewalks in the neighborhood where the retail outlet is located may increase the number of customers to visit that store.

Beyond the built environment, other political and economic conditions may affect sales of the item of interest. For fruits and vegetables, for example, there may also be a local or national campaign to increase fruit vegetable consumption that started during the intervention period. There may also have been a sugar sweetened beverage (SSB) tax passed during the time of your intervention which may reduce sales of SSBs, so the researcher should control for the introduction of those policies in regression models. An economic downturn or an increase in gas prices may contribute in a reduction of fruit and vegetables as a proportion of total sales as food budgets become more stretched. It is integral to think through what factors external to the store might affect sales prior to data collection and analysis so that these factors can be controlled for in analyses. Consultation with economists, public policy analysts, and geographers may be helpful in identifying key characteristics to monitor and include in analyses.

F. Loyalty Card Programs and Demographic Data

Loyalty card programs provide the opportunity to observe individual purchasing behavior. Here, sales data can come close to mimicking purchase data, which will improve the ability to make inferences at the individual level. However, not all retailers will have loyalty card programs, especially independents and/or smaller retail outlets. Because of the value of loyalty card program data, it is worthwhile to ask retail partners if they would be interested in starting a loyalty card program prior to data collection. Loyalty card programs may be another area where researchers can build in grant funding to offset the initial cost of loyalty card programs, providing further incentive for retailers to partner with researchers.

However, there are some concerns with working with loyalty card programs. For example, loyalty card customers may be more likely to respond to promotions, or may be of a different socioeconomic status compared to those who are not loyalty card customers. It will be difficult to compare the demographics of loyalty card users versus nonusers, so generalizability of analysis using loyalty card demographics is limited.

If researchers do want to use loyalty card programs, one option is to have a loyalty card drive or promotion prior to conducting the intervention in an attempt to increase the proportion of sales that are coming from loyalty card users. However, there still may be some selection bias in those who choose to sign up for a loyalty card program as opposed to those who do not want to sign up for loyalty card program.

It is important to note that researchers often cannot expect stores to provide customer demographic data alongside loyalty card sales data. Sharing these data can be explicitly forbidden as part of the loyalty card application signed by the customer. Furthermore, assuming demographic data can be shared, it is often sparse or misleading; customers are not required to provide accurate or up-to-date phone and home address information. Likewise, date of birth is not always required. Other demographic data, like gender and race, never exist because retailers do not ask for this information. Therefore, if selection bias is a concern, researchers would have to conduct demographic surveys of both loyalty and non-loyalty users to compare demographics of the two groups.

The value of loyalty card data, however, is not the demographic information, but the ability to link purchases to a single household or individual over time. Data that is linked to the same unit (household or person) over time is referred to as "Panel Data". Panel data is valuable because it expands the set of analytical tools available to the researcher.

G. **Identifying a Comparison Store**

If the researcher is implementing an in-store intervention, it is important to identify a control store or stores. The control store should have no similar intervention in-store throughout the research period. The control store allows the researcher to come as close as possible to the counterfactual, or what would have occurred in the intervention store if no intervention had occurred.

Ideally, the selection of stores as interventions or controls should be randomized so that confounders variables are distributed by chance between the two groups. Researchers may want to conduct a stratified random sample whereby retail outlets are randomly selected and then assigned within a particular category to ensure that intervention and control settings are comparable (i.e. randomization of just rural stores, or randomly selecting stores within urban and rural categories).

Retailers may be reluctant to participate in randomization because of the lack of benefit to the retailer of being in a control group. One potential solution is to conduct a delayed intervention whereby start dates of the intervention are randomized such that all retailers receive the intervention but the group(s) with the later start date(s) serve as the control group for the retailers with the earlier start dates. Another alternative is conducting a cross-over design study in which stores cross over from one treatment to another or from treatment to control during different time periods. We recommend a 30-day interval between start dates to allow for sufficient time for the intervention to affect customers. Balance checks for demographic and store characteristics should be conducted prior to analyzing data; these balance checks can be conducted using statistical software to determine if the intervention and control group are statistically significantly different based on demographic characteristics. The goal is to have control and intervention stores that are similar at baseline. If characteristics of intervention and control groups vary, then these characteristics should be controlled for in analyses.

If randomization of stores is not possible, then try and identify comparison stores that have similar characteristics as the intervention stores. Characteristics should be matched at the retailer level (e.g. size of retailer, urban/suburban/rural setting), as well as shopper characteristics. As described above, researchers may rely on census track data and/or store provided data to examine these characteristics.

H. **Empirical Methods and Analyses**

Data availability determines the set of available analytical tools that can be used to determine whether an intervention to promote healthier food choice was successful. The more disaggregated and detailed the data, the more powerful the analytical methods that can be brought to bear.

Ideally, the data provided by the retailer would be transaction-level data where each transaction could be linked to the same individual over time using a loyalty card number. Data with repeated observations from the same unit (store, person, household, etc.) over time is known as *panel data*. Transaction-level data from the same individual over time would be *individual-level* panel data.

Panel data is powerful because it allows the researcher to control for unobserved effects at every known level of analysis. For example, assume the researcher has transaction-level data where purchases are linked to individuals over multiple months. We can safely assume purchases can be linked to stores and that we know the date and time of purchase. With these panel data, the researcher can run regression models with *individual and store fixed effects*.

Fixed effects models are powerful because the researcher can control for any *static* variables (variables that do not change over time) along identifiable levels. For example, at the individual level, the researcher can control for static variables like gender and race. At the store level, the researcher can control for the effect of the store location and size. Even unmeasurable static variables, like store reputation or individual food preferences (vegan or vegetarian), are captured by fixed effects.

The reverse is also true. If loyalty card data is unavailable, the researcher loses the ability to capture individual fixed effects. However, if the researcher can still link purchases to stores via obtaining sales data, then it would still be possible to run models with store fixed effects. This is an example of why more detail is always preferable to less.

Transaction-level data is valuable even without loyalty card data. The researcher can observe how products are purchased together ("bundled"). Bundling helps determine complementary goods and can also help determine any unintended consequences of changing relative prices. For example, "self-licensing" could be triggered by making healthy products cheaper. In response to buying a discounted healthy snack (good behavior), a customer may "self-license" to also buy an unhealthy snack.

Furthermore, estimating average treatment effects remains equally plausible when the researcher can aggregate the transaction data as needed. For example, daily total sales data by product is useful for making distributional comparisons between two experimental groups. Daily total sales data by product can be easily generated from transaction-level data, but also obtained without data being available at the transaction-level.

If data from comparison stores are available, the researcher can expand to using models such as Difference-in-Differences (DD). DD models help the researcher control for any dynamic effects shared by both the experimental and comparison stores. For example, overall increases in sales of a particular product due to a promotion unrelated to your intervention could be "differenced out". Likewise, DD can control for macroeconomic conditions affecting all stores such as recessions.

DD models are still powerful even without access to transactional-level panel data. For example, if the research question of interest compares the change in sales between groups of stores for a specific product over time, then repeated cross sections (RCS) data is sufficient *as long as the experimental and control stores can be identified*. Transactions without loyalty card data or aggregated product sales data at the store level would be considered RCS data. DD is still feasible if the experiment and comparison stores can be identified and data is available before, during, and after the start of the experiment.

The following graphic provides a visual for the different types of data and analytical procedures discussed above. The top of the inverted pyramid summarizes the most ideal type of data and analytic procedures; as the reader moves down the pyramid, the types of data and analyses become less ideal.

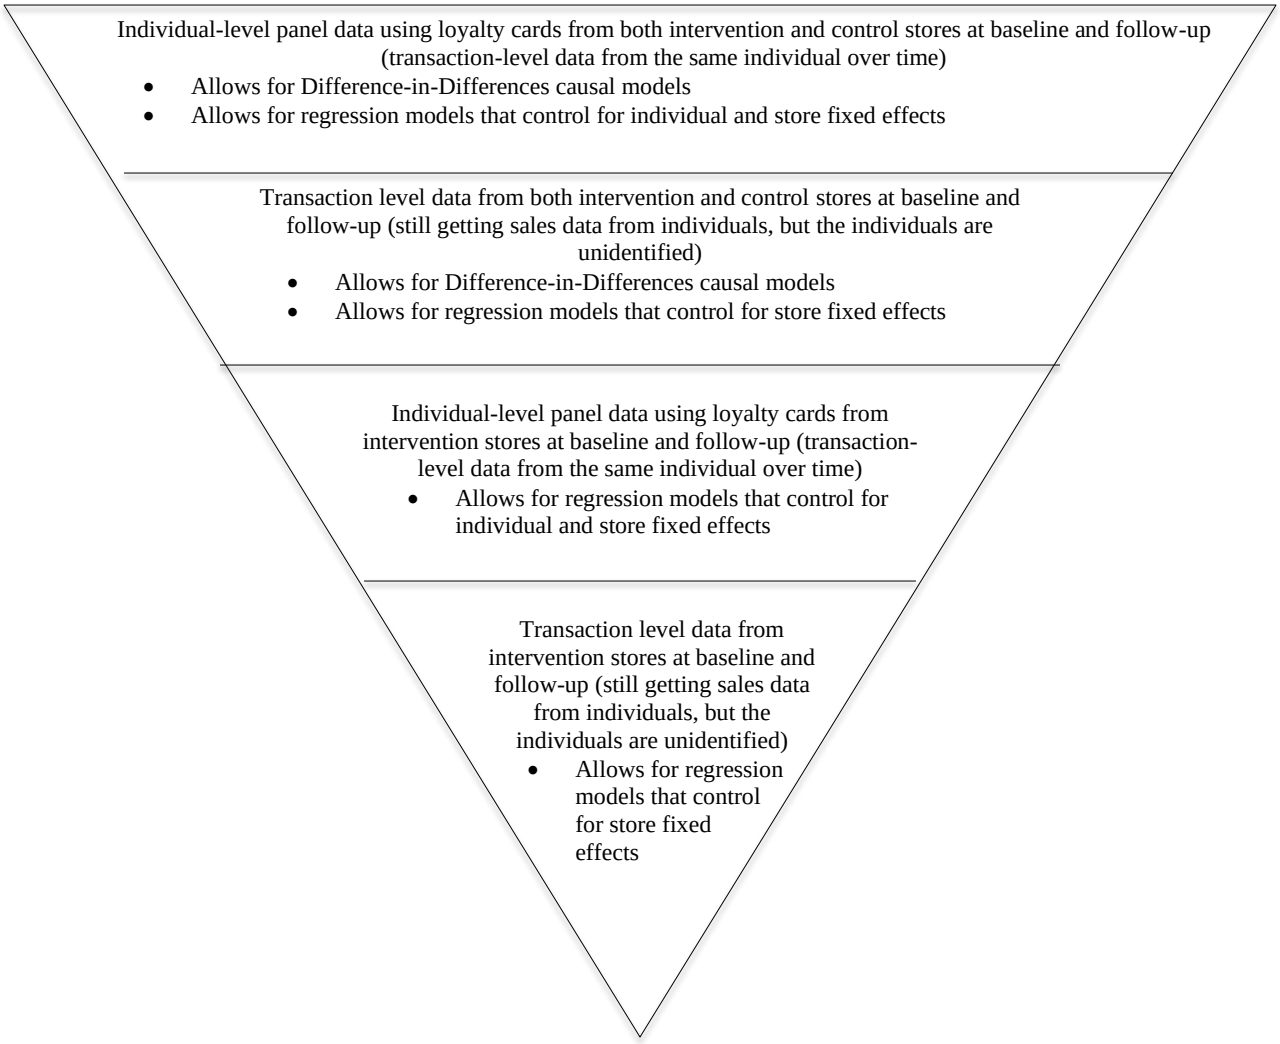

In economics, the standard is to perform "robustness checks". Robustness checks refers to the process of running the same basic model but with varying controls. In the case of a DD model, the "basic" model would be the model without any covariates. The researcher creates a "regression table" of the coefficient estimates, focusing on the coefficient of interest (e.g. change in banana sales due to treatment). The researcher then runs a variation on the basic model by adding a set of covariates (e.g. controls for seasonality). These coefficient estimates are also added to the regression table. This process repeats for every variation of controls. The aim is to observe how the coefficient of interest (e.g. changes in banana sales) responds to changing model specifications. A "robust" coefficient would be one that does not vary dramatically in response to covariate changes. This process of adding and removing covariates, however, requires the researcher to carefully consider the relationship between the dependent variable, the coefficient of interest, and important covariates.

For more on panel data methods and difference-in-differences, see Wooldridge (2010)[10] and Angrist and Pischke (2008).[11] For details on robustness checks, see Lu and White (2014).[12]

Areas for Future Research

The USDA-SR contains thousands of branded food products. However, there is a need for additional research and development for smartphone apps and software to researchers to more easily map sales data onto USDA-SR items and food groups.

Websites that Provide Additional Resources about Working with Sales Data

Food Marketing Institute (fmi.org/research-resources): The Food Marketing Institute conducts research on a variety of topics related to grocery retailers, including sales trends and external factors that affect food retail sales.

National Bureau of Economic Research (nber.org): The National Bureau of Economic Research publishes white papers on a variety of economic issues, including relationships between food prices, economic indicators, and sales.

The Healthy Food Access Portal (healthyfoodaccess.org): The Healthy Food Access Portal is a project of PolicyLink, the Food Trust, and Reinvestment Fund, and provides resources for the implementation of interventions and analysis of retail data.

The Duke-UNC Center for Behavioral Economics and Healthy Food Choice Research (becr.sanford.duke.edu): Alongside this brief, several other briefs are available on the application of behavioral economics to retail setting to encourage healthy food choices, particularly among federal nutrition program recipients.

Conclusions

Procuring and analyzing sales data is often challenging, but using sales data as an outcome metric, rather than self-report measures, can greatly improve a study's internal validity due to the objectivity of the data. Sales data changes are also easily translatable for policy and other non-research audiences. Best practices for analyzing sales data to assess the effectiveness of healthy retail interventions are in the early stages of development. This brief aims to provide a starting point for intervention researchers interested in procuring, storing, and analyzing sales data.

Sources:

1. Escaron AL, Meinen AM, Nitzke SA, Martinez-Donate AP. Supermarket and Grocery Store–Based Interventions to Promote Healthful Food Choices and Eating Practices: A Systematic Review. *Prev Chronic Dis.* 2013;10:120156.

2. Donaldson, S.I. & Grant-Vallone, E.J. Understanding Self-Report Bias in Organizational Behavior Research. *Journal of Business and Psychology.* 2002;17: 245.

3. Gittelsohn J, Laska MN, Karpyn A, Klingler K, Ayala GX. Lessons Learned From Small Store Programs to Increase Healthy Food Access. *American Journal of Health Behavior.* 2014;38(2):307-315.

4. De Marco, M, Chapman, L, Siddique, N. What's in it for Retailers? Establishing Partnerships with Food Retailers to Conduct Healthy Food Choice Research. Available at: https://becr.sanford.duke.edu/wp-content/uploads/2016/12/PartneringWithRetailers_Final.pdf. Published December 2016. Accessed February 19th, 2017.

5. Healthy Food Retail Working Group. *Grocery Retailer Academic Collaborative Guidelines (GRAC).* Healthy Eating Research. Nutrition and Obesity Policy Research and Evaluation. 2017. Available at: https://nopren.org/wp-content/uploads/2017/05/GRAC-Guidelines.pdf. Accessed June 5, 2017.

6. The U.S. Department of Justice. Guidelines for a Memorandum of Understanding. http://www.doj.state.or.us/victims/pdf/mou_sample_guidelines.pdf. Accessed February 19th, 2017.

7. United States Department of Agriculture. Agriculture Research Services. Food Composition Database. Available at: https://ndb.nal.usda.gov/ndb/. Accessed on March 3, 2017.

8.Brinkerhoff, K. M., Brewster, P. J., Clark, E. B., Jordan, K. C., Cummins, M. R., & Hurdle, J. F. (2011). Linking supermarket sales data to nutritional information: an informatics feasibility study. *AMIA Annual Symposium Proceedings. AMIA Symposium*, *2011*, 598–606. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/22195115. Accessed on March 3, 2017.

9. Pennington JA, Stumbo PJ, Murphy SP, McNutt SW, Eldridge AL, McCabe-Sellers BJ, Chenard CA. Food composition data: the foundation of dietetic practice and research. J Am Diet Assoc. 2007 Dec;107(12): 2105-

10. Wooldridge, Jeffrey M. *Econometric Analysis of Cross Section and Panel Data*. MIT Press, 2010.

11. Angrist, Joshua D., and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton university press, 2008.

12. Lu, Xun, and Halbert White. "Robustness Checks and Robustness Tests in Applied Economics." *Journal of Econometrics*, Annals Issue: Misspecification Test Methods in Econometrics, 178, Part 1 (January 2014): 194–206. doi:10.1016/j.jeconom.2013.08.016.