Trends of Tobacco Compliance in Wyoming

Lucy Ward

University of Wyoming Honors Capstone

Mentor: Dr. Annalisa Piccorelli

Introduction

The purpose of the annual Synar Tobacco inspections is to determine the level of tobacco compliance in the state of Wyoming in order to evaluate prevention needs. This research using eight years of Synar data seeks to explore the factors that make a location more likely to sell a tobacco product to an underage buyer. During each inspection, the youth inspector makes observations about the store they enter and the interaction they have with the clerk, and then record those observations in a detailed form immediately after leaving the store. We therefore have data regarding the youth inspector's gender and age, what type of store location they are entering (gas, grocery, or other for these purposes), whether products were readily accessible for customers to pick up or whether they had to ask a clerk to access the products, whether youth access signs (WE ID, etc.) were present in the store, the perceived gender and age of the clerk they interacted with, the type of tobacco they requested (cigarettes or chewing tobacco), whether the clerk asked for their ID, and whether the clerk asked for their age. Summaries of these outcome variables can be seen in Table 1.

Table 1: Frequencies and Percentages of Predictor Variables

Predictor Variable	Variable Value	Frequency (Percentage)
Year	2012	296 (14.4%)
	2013	368 (17.9%)
	2014	242 (11.8%)
	2015	230 (11.2%)
	2016	177 (8.6%)
	2017	289 (14.1%)
	2018	256 (12.5%)
	2019	198 (9.6%)
Inspector Gender	Female	1047 (51.0%)
	Male	1007 (49.0%)
Inspector Age	16	865 (42.1%)
	17	1190 (57.9%)
Store Type	Gas Station	1228 (62.7%
	Grocery Store	436 (21.2%)
	Other	331 (16.1%)
Product Accessibility	Accessible	49 (2.4%)
	Not Accessible	1984 (97.6%)
Youth Access Signs	Present	1765 (86.9%)
	Not Present	267 (13.1%)
Clerk Gender	Female	1531 (74.7%)
	Male	520 (25.4%)
Tobacco Type	Cigarettes	1467 (71.4%)
	Chewing Tobacco	588 (28.61%)
Did Clerk Ask for ID?	Yes	1847 (90.5%)
	No	196 (9.6%)
Did Clerk Ask for Age?	Yes	159 (7.4%)
	No	1865 (92.6%)
Clerk Age	Mean (SD)	38.8 (14.5)

Before beginning analysis, it was predicted that the clerk asking the inspector for their ID would reduce the likelihood of a successful buy attempt. It was also predicted that older inspectors would be more likely to successfully purchase the products. An expectation was that inspectors purchasing cigarettes would be more successful in their purchases than those requesting chewing tobacco.

Another curiosity was the presence of potential interactions between clerk and inspector gender, in which it was predicted that when the clerk and inspector have the same gender a successful buy attempt would be more likely.

Methods

Each summer, the Wyoming Survey and Analysis Center (WYSAC) executes Synar Tobacco Inspections on behalf of the Wyoming Department of Health. A random sample of tobacco retailers was generated using the Synar Survey Estimation System (SSES) sampling tool. Youth inspectors aged 16 and 17 were recruited from across the state to complete the inspections (n = 2056). Drivers based in Laramie were hired to drive youth inspectors to sampled location, train the inspectors in Synar procedures, and ensure proper collection of data.

When a youth inspector entered a location, they approached the clerk and requested a tobacco product (either chewing tobacco or cigarettes based on instruction from the driver). If the clerk asked the inspector for ID and refused to complete the sale when the inspector failed to produce ID, that was considered an unsuccessful buy attempt, a "pass" for the retailer. However, if the clerk rang up the purchase and was prepared to make the sale, the inspector "realized" they did not have enough money to purchase the product and left the retailer. This was considered a successful buy attempt, or a "fail" for the retailer. After leaving the location, the inspector and the driver worked together to record the relevant information from the inspection based on the inspector's interaction with the clerk and other observations in the store. Prior to 2019, data was recorded on paper then manually entered into the computer. In 2019, the WYSAC System Administrator developed a mobile app that allowed inspectors and drivers to enter data from the inspection immediately upon its completion. In 2020, due to the passage of Tobacco 21, these procedures were adjusted to allow for 18 to 20-year-old inspectors. For this reason, data from 2020 were excluded from this research.

Exploration and analysis of the data were done using SAS software. Summary measures regarding each variable's distribution relative to inspection outcome were computed and Chi-Square tests for binary predictor variables and a t-test for a continuous predictor variable to determine significant differences across inspection outcomes were computed.

Using the logistic procedure, a full model including all linear predictors as well as interactions between inspector gender and age, clerk and inspector gender, and inspector gender and tobacco type was constructed. Backward elimination was used to remove one predictor at a time, and the final model was selected based on optimized AIC and concordance. In the final model, some predictors are not statistically significant, but this was the candidate model that simultaneously maximized concordance and minimized AIC relative to the others.

In order to make observations about the effects of different predictors over time, the data were stratified by year and simplified logistic regression models were constructed to relate each predictor with the inspection outcome for each year. The resulting coefficients and standard error values were used to make observations about changes in the effects of different predictors over time.

Results

The distribution of each predictor variable by outcome can be seen in Table 2. Significant differences in outcome were observed for year, inspector age, clerk asking for ID, and clerk age.

Table 2: Frequencies and Percentages of Predictor Variables by Outcome

Inspection Outcome		Buy Attempt Successful 166 (8.1%)	Buy Attempt Unsuccessful 1890 (91.9%)
Year*	2012	43 (25.1%)	253 (13.4%)
	2013	32 (19.3%)	336 (17.8%)
	2014	19 (11.5%)	223 (11.8%)
	2015	12 (7.2%)	218 (11.5%)
	2016	21 (12.7%)	156 (8.3%)
	2017	15 (9.0%)	274 (14.5%)
	2018	15 (9.0%)	241 (12.8%)
	2019	9 (5.4%)	189 (10.0%)
Inspector Gender	Female	95 (57.2%)	952 (50.4%)
	Male	71 (42.8%)	936 (49.6%)
Inspector Age*	16	56 (33.7%)	809 (42.8%)
	17	110 (66.3%)	1080 (57.1%)
Store Type	Gas Station	100 (60.2%)	1188 (62.9%)
	Grocery Store	32 (19.3%)	404 (21.4%)
	Other	34 (20.5%)	298 (15.8%)
Product Accessibility	Accessible	3 (1.8%)	46 (2.5%)
	Not Accessible	163 (98.2%)	1821 (97.5%)
Youth Access Signs	Present	135 (83.9%)	1630 (87.1%)
	Not Present	26 (16.2%)	241 (12.9%)
Clerk Gender	Female	114 (68.7%)	1417 (75.2%)
	Male	52 (31.3%)	468 (24.8%)
Tobacco Type	Cigarettes	109 (65.7%)	1358 (71.9%)
	Chewing Tobacco	57 (34.3%)	531 (28.1%)
Did Clerk Ask for ID?*	Yes	36 (21.7%)	1811 (96.5%)
	No	130 (78.3%)	65 (3.5%)
Did Clerk Ask for Age?	Yes	13 (7.9%)	136 (7.4%)
	No	151 (92.1%)	1714 (92.7%)
Clerk Age*	Mean (SD)	34.8 (14.5)	39.2 (14.4)

*indicates a significant difference between groups (p-value < 0.05)

The full model with all linear predictors as well as three interaction terms had AIC of 534.303 and concordance of 0.892. After the initial model, one covariate at a time was removed based on its p-value. The terms removed, in sequence, were a store type of gas station, youth access signs, the interaction between clerk and inspector gender, the interaction between inspector age and inspector gender, the interaction between inspector gender and tobacco type, and tobacco type.

The final model had predictors year, inspector gender, inspector age, store type (grocery store or other), whether tobacco products were readily accessible to customers, clerk gender, clerk age, whether the clerk asked for id, and whether the clerk asked for gender. As stated, this model was selected because it minimized AIC and maximized concordance relative to other candidate models. Estimated coefficients, p-values, estimated odds ratios, and estimated relative risks for this model can be seen in Table 3.

Table 3: Final Model Output

Predictor Variable	Variable Level	Maximum Likelihood Estimate	Wald Chi-Square p-value	Estimated Odds Ratio[+]	Estimated Relative Risk[++]
Intercept		-2.3308	< 0.0001*		
Year	2012 vs. 2019	0.8936	0.0014*	4.409	2.3052
	2013 vs. 2019	0.0369	0.8972	1.872	1.0369
	2014 vs. 2019	-0.2958	0.4093	1.342	0.7528
	2015 vs. 2019	0.1418	0.7197	2.079	1.1460
	2016 vs. 2019	0.3410	0.3709	2.537	1.3838
	2017 vs. 2019	-0.5894	0.1256	1.001	0.0568
	2018 vs. 2019	0.0619	0.8661	1.919	1.0620
Inspector Gender	Female vs. Male	0.1385	0.2841	1.319	1.1424
Inspector Age	16 vs. 17	-0.2045	0.1392	0.664	0.8222
Store Type	Grocery Store vs. Other	0.2551	0.1223	1.666	1.2761
Product Accessibility	Not Accessible vs. Accessible	0.4646	0.3230	2.533	0.1554
Clerk Gender	Female vs. Male	-0.1802	0.1939	0.697	0.8417
Did Clerk Ask for ID?	No vs. Yes	2.7323	< 0.0001*	236.158	9.5641
Did Clerk Ask for Age?	No vs. Yes	1.2165	< 0.0001*	11.393	3.0697
Clerk Age	1-year increase	-0.0197	0.0322*	0.980	0.9813

*indicates a significant difference in outcome between the two variable levels (p-value < 0.05)

+while all other variables are held constant

++calculated using baseline values of 2019, male inspector, 17-year-old inspector, other store type, products accessible, male clerk, did ask for ID, did ask for age, and clerk age 40

A Hosmer-Lemeshow Goodness-of-Fit test showed that this model did not provide a good fit to the data. However, “several studies have documented the frequent rejection of acceptable – though nonperfect – models” (Nattino, Pennell, and Lemeshow, 2020). Thus, the results of this model were still used to make some inferences. Evaluation of model residuals and influence plots showed that outliers may have caused the lack of fit. Because of confidence that these outliers were legitimate data points, they were not removed from consideration.

Significant predictors in the final model were the year 2012, clerk age, whether the clerk asked for ID, and whether the clerk asked for age. The estimated odds of a successful buy attempt in 2012 were 4.409 times, or 340.6% higher than, that of 2019 when all other variables were held constant. Coefficients for other years were not significant. For clerk age, the estimated odds of a successful buy attempt decreased by 0.980 times, or 2% on average, as the estimated age of the clerk increased by one year while all other variables were held constant. This suggests that older clerks were less willing to sell to underage buyers. When a clerk failed to ask for the youth inspector’s ID, the estimated odds of a successful buy attempt were 235.150 times higher than that of an inspection where the clerk asked for ID when all other variables were held constant. Additionally, when a clerk failed to ask for an inspector’s age the odds of a successful buy attempt were 1,039.3% higher than that of an inspection where the clerk did ask for age where all other variables were held constant. Asking for ID and age were clear and strong predictors of inspection outcome.

Tobacco compliance is changing in Wyoming. A Cochran-Mantel-Haenszel test, which tests for relationships involving ordinal variables, showed that there is a linear association

between year and inspection outcome. This is reinforced by Figure 1, which shows the inspection failure rates for each year as well as a trend line that summarizes the rates over time. In general, inspection failure rates decreased across the years included in these data. Lower rates of inspection failures mean that fewer underage buyers are gaining access to tobacco products, so this is good news for Wyoming.

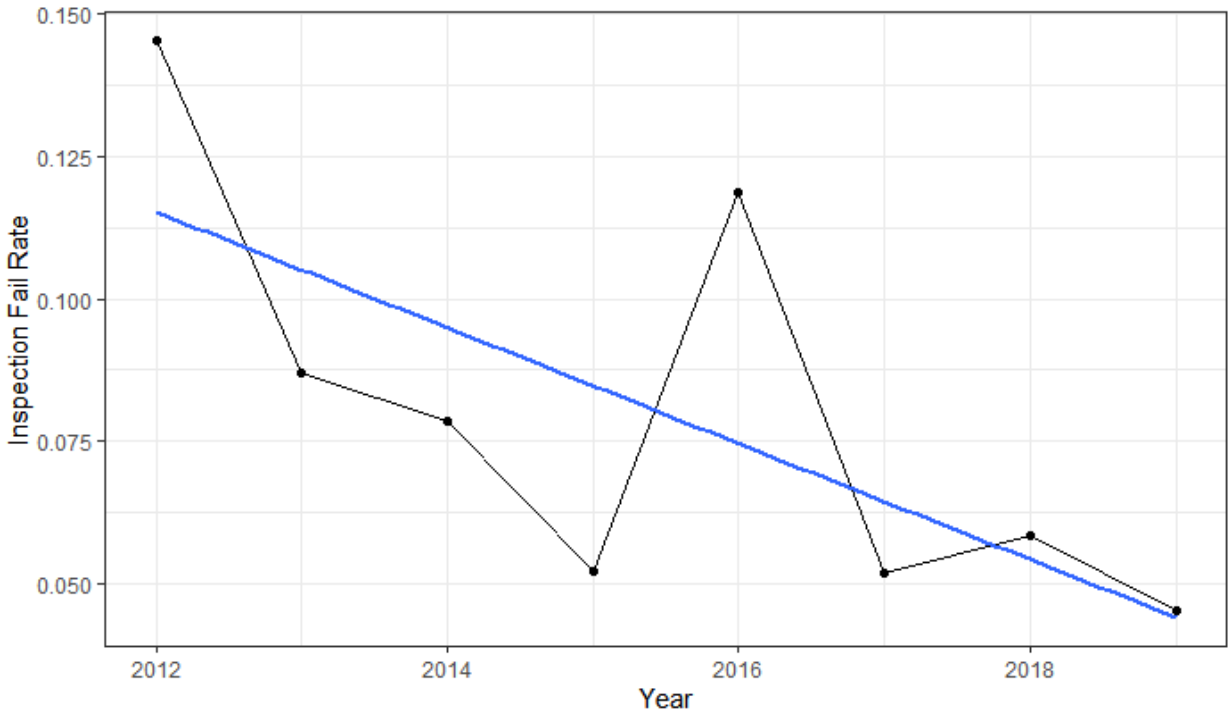

Figure 1: Inspection Failure Rates are Generally Decreasing

Looking at the results of the stratification, some interesting patterns emerged. Figure 2 shows the coefficient estimates for the effect of asking for ID over time, as well as their standard errors and their proximity to zero. This plot shows that there were not major differences in the effect across time, because most of the error bars for the estimates overlapped with one another. However, none of the error bars overlapped with zero, or even came close, which showed a strong and consistent effect of this predictor.

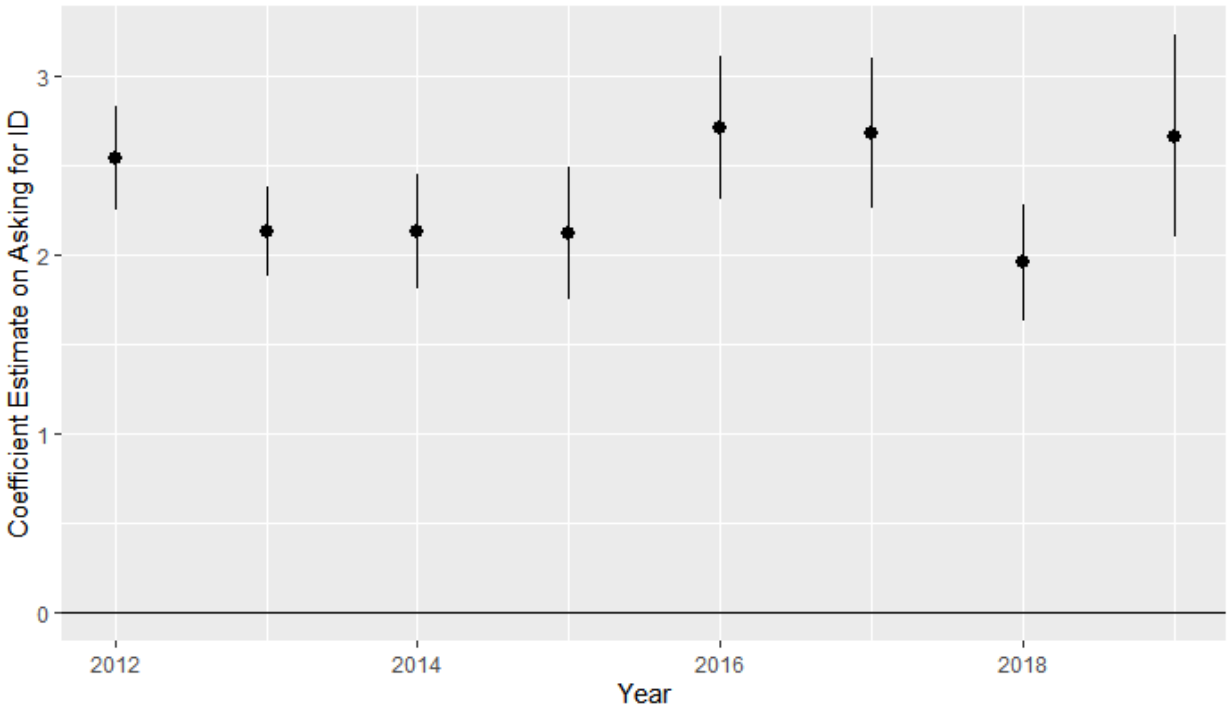

Figure 2: Effect of Asking for ID Over Time (asking for ID vs. not asking for ID)

Figure 3 illustrates the effect of inspector gender across the observed years of inspections. In some years, when the error bar intersected with zero, there was little effect of gender on inspection outcome. In 2013, 2015, and 2016, when the coefficient value was positive, female inspectors were more likely to have a successful buy attempt. In 2018, male inspectors were more likely to have a successful buy attempt. There was variation in the effect of inspector gender over time, but when the data were combined in a full model, the inspector being female increased the likelihood of a successful buy attempt.

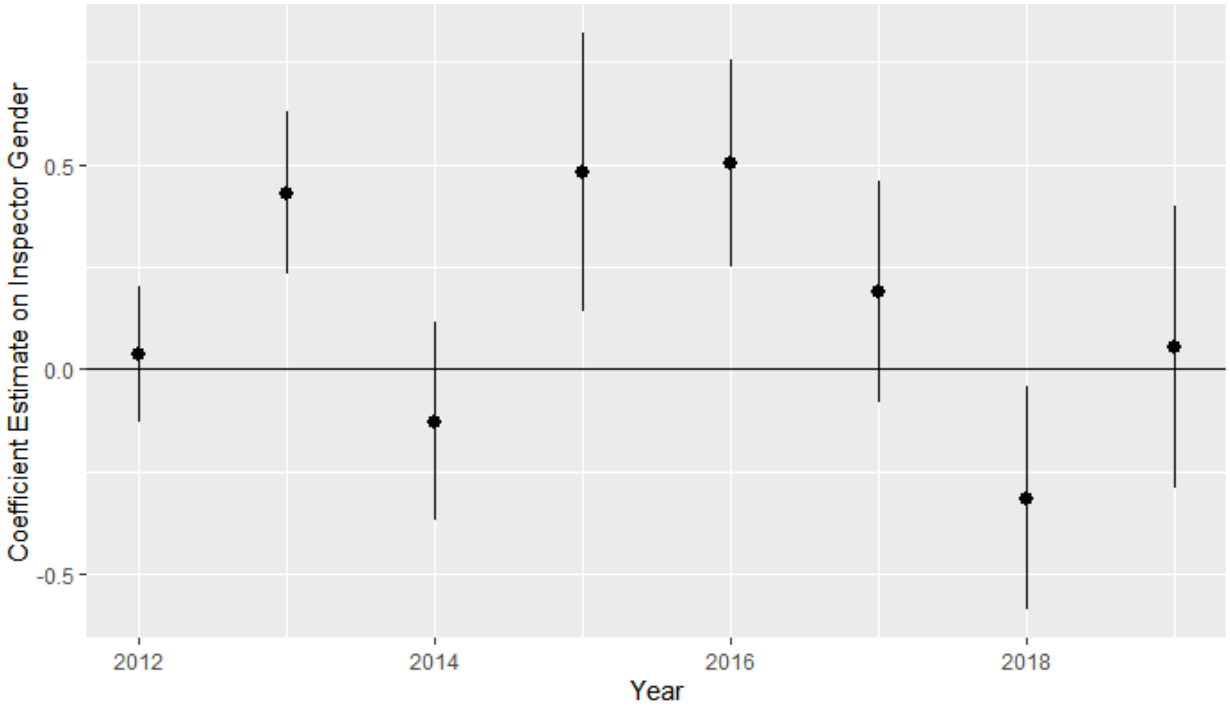

Figure 3: Effect of Inspector Gender Over Time (female vs. male)

Finally, Figure 4 shows this same statistic for clerk gender. Most years had little effect of this predictor on inspection outcome, but 2015, 2018, and 2019 found that male clerks were more likely to allow sales for underage buyers. The error bars for most years overlapped, suggesting there was not much change in this effect from 2012 to 2018, but 2019 appeared to have had a different coefficient value from 2012 and 2013.

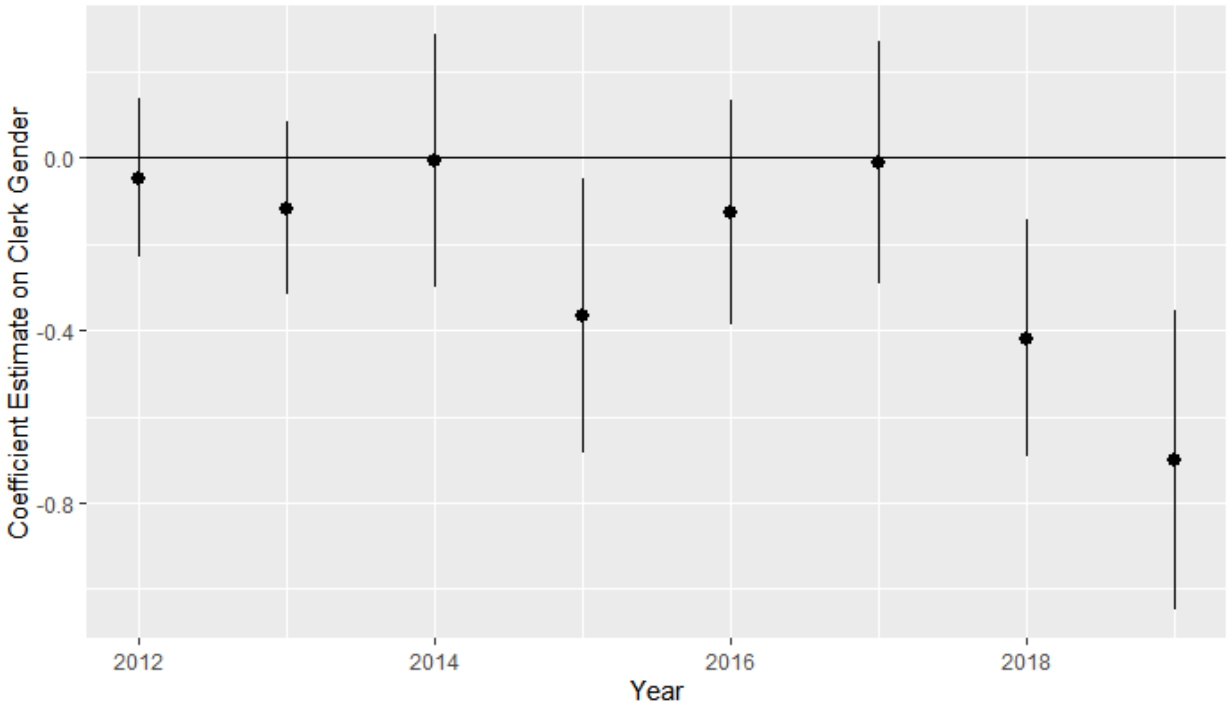

Figure 4: Effect of Clerk Gender Over Time (female vs. male)

Discussion

Based on the findings of the logistic regression model, best practices for preventing sale to underage individuals is to train clerks to consistently ask for identification and/or ask for age. Inspection failure rates are trending downward, which means that prevention efforts and policies over the past eight years have helped to reduce sales to underage consumers.

The strengths of this study were its large sample size of over 2,000 inspections across eight years and data quality from consistent reporting and formatting methods. This large dataset provides an opportunity to find patterns of tobacco compliance that may not emerge during a single year of inspection. One potential challenge is that since clerk asking for ID is such a strong indicator (as is expected) that other factors' effects were minimized. Future exploration focused on other predictors like product accessibility and youth access signs could provide insight to inform policy or best practices of such measures.

Development of a mobile application to collect Synar inspection data has allowed for collection of precise time and location data for each inspection. As this app continues to be used, patterns across time, both time of day and time of year, and space could be explored and emerged. Spatial patterns could reveal certain communities or counties struggling with tobacco compliance and help to inform future prevention measures.

In the year 2020, new legislation made the age to purchase tobacco products 21 instead of 18. Inspection protocols were updated to use 18- to 20-year-old inspectors, which will likely continue in future years. Data from 2020 was excluded from this project because of this change,

but after more data is collected with older inspectors it will be interesting to see if and how these patterns change.

References

Nattino, G, Pennell, ML, Lemeshow, S. Assessing the goodness of fit of logistic regression models in large samples: A modification of the Hosmer-Lemeshow test. *Biometrics*. 2020; 76: 549– 560. https://doi.org/10.1111/biom.13249

Appendix

Coding in SAS for Analysis

```
*import data from Excel file;
PROC IMPORT OUT= WORK.dataclean
            DATAFILE=
"\\dellnas\lward7\Folders\Desktop\Thesis\CompiledCleaned.
xlsx"
            DBMS=EXCEL REPLACE;
     RANGE="Sheet1$";
     GETNAMES=YES;
     MIXED=NO;
     SCANTEXT=YES;
     USEDATE=YES;
     SCANTIME=YES;
RUN;

*obtain summary measures of predictor variables;
proc freq data = work.dataclean;
tables year inspector_gender inspector_age store_type_gas store_type_grocery
product_accessibility youth_access_signs clerk_gender tobacco_type ask_for_id
ask_for_age buy_attempt_successful;
run;

proc means data = work.dataclean;
var clerk_age;
run;

*obtain summary measures of predictor variables by the outcome of the
inspection;
proc freq data = work.dataclean;
tables (year inspector_gender inspector_age store_type_gas store_type_grocery
product_accessibility youth_access_signs clerk_gender tobacco_type ask_for_id
ask_for_age) * (buy_attempt_successful) / chisq measures;
run;

proc means data = work.dataclean;
by buy_attempt_successful;
var clerk_age;
run;

proc ttest data = work.dataclean;
class buy_attempt_successful;
var clerk_age;
run;
```

```
*Cochran-Mantel-Haenszel test for linear association between year and
inspection outcome;
proc freq data = work.dataclean;
tables year*buy_attempt_successful / cmh;
run;

*construct a full logistic model including all linear predictors;
proc logistic data = work.dataclean descending;
class year inspector_gender inspector_age store_type_gas store_type_grocery
product_accessibility youth_access_signs clerk_gender tobacco_type ask_for_id
ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_gas store_type_grocery product_accessibility youth_access_signs
clerk_gender clerk_age tobacco_type ask_for_id ask_for_age;
run;

*construct secondary full model including potential interactions;
proc logistic data = work.dataclean descending;
class year inspector_gender inspector_age store_type_gas store_type_grocery
product_accessibility youth_access_signs clerk_gender tobacco_type ask_for_id
ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_gas store_type_grocery product_accessibility youth_access_signs
clerk_gender clerk_age tobacco_type ask_for_id ask_for_age
inspector_gender*inspector_age clerk_gender*inspector_gender
inspector_gender*tobacco_type;
run;

*remove one term iteratively based on p-value;
proc logistic data = work.dataclean descending;
class year inspector_gender inspector_age store_type_grocery
product_accessibility youth_access_signs clerk_gender tobacco_type ask_for_id
ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_grocery product_accessibility youth_access_signs clerk_gender
clerk_age tobacco_type ask_for_id ask_for_age inspector_gender*inspector_age
clerk_gender*inspector_gender inspector_gender*tobacco_type;
run;

proc logistic data = work.dataclean descending;
class year inspector_gender inspector_age store_type_grocery
product_accessibility clerk_gender tobacco_type ask_for_id ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_grocery product_accessibility clerk_gender clerk_age tobacco_type
ask_for_id ask_for_age inspector_gender*inspector_age
clerk_gender*inspector_gender inspector_gender*tobacco_type;
run;

proc logistic data = work.dataclean descending;
class year inspector_gender inspector_age store_type_grocery
product_accessibility clerk_gender tobacco_type ask_for_id ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_grocery product_accessibility clerk_gender clerk_age tobacco_type
ask_for_id ask_for_age inspector_gender*inspector_age
inspector_gender*tobacco_type;
run;
```

```
proc logistic data = work.dataclean descending;
class year inspector_gender inspector_age store_type_grocery
product_accessibility clerk_gender tobacco_type ask_for_id ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_grocery product_accessibility clerk_gender clerk_age tobacco_type
ask_for_id ask_for_age inspector_gender*tobacco_type;
run;

proc logistic data = work.dataclean descending;
class year inspector_gender inspector_age store_type_grocery
product_accessibility clerk_gender ask_for_id ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_grocery product_accessibility clerk_gender clerk_age ask_for_id
ask_for_age;
run;

proc logistic data = work.dataclean descending;
class year inspector_gender inspector_age store_type_grocery clerk_gender
ask_for_id ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_grocery clerk_gender clerk_age ask_for_id ask_for_age;
run;

proc logistic data = work.dataclean descending;
class year inspector_age store_type_grocery clerk_gender ask_for_id
ask_for_age;
model buy_attempt_successful = year inspector_age store_type_grocery
clerk_gender clerk_age ask_for_id ask_for_age;
run;

proc logistic data = work.dataclean descending;
class year inspector_age store_type_grocery ask_for_id ask_for_age;
model buy_attempt_successful = year inspector_age store_type_grocery
clerk_age ask_for_id ask_for_age;
run;

proc logistic data = work.dataclean descending;
class year store_type_grocery ask_for_id ask_for_age;
model buy_attempt_successful = year store_type_grocery clerk_age ask_for_id
ask_for_age;
run;

proc logistic data = work.dataclean descending;
class year ask_for_id ask_for_age;
model buy_attempt_successful = year clerk_age ask_for_id ask_for_age;
run;

*select final model based on AIC and concordance;
proc logistic data = work.dataclean descending;
class year inspector_gender inspector_age store_type_grocery
product_accessibility clerk_gender ask_for_id ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_grocery product_accessibility clerk_gender clerk_age ask_for_id
ask_for_age;
run;
```

```
*evaluate Pearson deviance and goodness of fit statistics on final model;
proc logistic data = work.dataclean descending;
class year inspector_gender inspector_age store_type_grocery
product_accessibility clerk_gender ask_for_id ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_grocery product_accessibility clerk_gender clerk_age ask_for_id
ask_for_age / aggregate scale = none;
run;

*evaluate Hosmer-Lemeshow test on final model;
proc logistic data = work.dataclean descending;
class year inspector_gender inspector_age store_type_grocery
product_accessibility clerk_gender ask_for_id ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_grocery product_accessibility clerk_gender clerk_age ask_for_id
ask_for_age / lackfit;
run;

*evaluate influence diagnostics of final model;
proc logistic data = work.dataclean descending;
class year inspector_gender inspector_age store_type_grocery
product_accessibility clerk_gender ask_for_id ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_grocery product_accessibility clerk_gender clerk_age ask_for_id
ask_for_age / influence;
run;

*compute streschi residuals of final model;
proc genmod data=work.dataclean descending;
class year inspector_gender inspector_age store_type_grocery
product_accessibility clerk_gender ask_for_id ask_for_age;
model buy_attempt_successful = year inspector_gender inspector_age
store_type_grocery product_accessibility clerk_gender clerk_age ask_for_id
ask_for_age / dist=bin link=log obstats residuals;
ods output obstats=resids;
run;

proc print data=resids;
var buy_attempt_successful year inspector_gender inspector_age
store_type_grocery product_accessibility clerk_gender clerk_age ask_for_id
ask_for_age streschi;
run;


*sort the data by year;
proc sort data = work.dataclean;
by year;
run;
```

```
*stratify the final model by year;
proc logistic data = work.dataclean descending;
by year;
class buy_attempt_successful inspector_gender inspector_age
store_type_grocery product_accessibility clerk_gender ask_for_id ask_for_age;
model buy_attempt_successful = inspector_gender inspector_age
store_type_grocery product_accessibility clerk_gender clerk_age ask_for_id
ask_for_age;
run;

*stratify the effect of inspector gender by year;
proc logistic data = work.dataclean descending;
by year;
class inspector_gender;
model buy_attempt_successful = inspector_gender;
run;

*stratify the effect of inspector age by year;
proc logistic data = work.dataclean descending;
by year;
class inspector_age;
model buy_attempt_successful = inspector_age;
run;

*stratify the effect of store type (grocery) by year;
proc logistic data = work.dataclean descending;
by year;
class store_type_grocery;
model buy_attempt_successful = store_type_grocery;
run;

*stratify the effect of product accessibility by year;
proc logistic data = work.dataclean descending;
by year;
class product_accessibility;
model buy_attempt_successful = product_accessibility;
run;

*stratify the effect of clerk gender by year;
proc logistic data = work.dataclean descending;
by year;
class clerk_gender;
model buy_attempt_successful = clerk_gender;
run;

*stratify the effect of asking for id by year;
proc logistic data = work.dataclean descending;
by year;
class ask_for_id;
model buy_attempt_successful = ask_for_id;
run;
```

```
*stratify the effect of asking for age by year;
proc logistic data = work.dataclean descending;
by year;
class ask_for_age;
model buy_attempt_successful = ask_for_age;
run;

*stratify the effect of clerk age by year;
proc logistic data = work.dataclean descending;
by year;
model buy_attempt_successful = clerk_age;
run;
```

Coding in R for Visualization

```
# read in the data
library(readxl)
by_year <- read_excel("by year.xlsx")
View(by_year)

# plot inspection failure rate by year
by_year %>%
  ggplot(aes(x = year,
             y = failrate)) +
  geom_line() +
  geom_point() +
  stat_smooth(method = lm,
              se = FALSE) +
  labs(x = "Year",
       y = "Inspection Fail Rate",
       title = "Figure 1: Inspection Failure Rates are Generally Decreasing")
+
  theme_bw()

# plot inspector gender coefficient by year
by_year %>%
  ggplot(aes(x = year,
             y = inspgenderest)) +
  geom_point() +
  geom_pointrange(aes(ymax = inspgenderest + inspgenderse,
                      ymin = inspgenderest - inspgenderse)) +
  geom_hline(aes(yintercept = 0)) +
  labs(x = "Year",
       y = "Coefficient Estimate on Inspector Gender",
       title = "Figure 3: Effect of Inspector Gender Over Time")
```

```r
# plot inspector age coefficient by year
by_year %>%
  ggplot(aes(x = year,
             y = inspageest)) +
  geom_point() +
  geom_pointrange(aes(ymax = inspageest + inspagese,
                      ymin = inspageest - inspagese)) +
  geom_hline(aes(yintercept = 0))

# plot store type coefficient by year
by_year %>%
  ggplot(aes(x = year,
             y = storegroceryest)) +
  geom_point() +
  geom_pointrange(aes(ymax = storegroceryest + storegroceryse,
                      ymin = storegroceryest - storegroceryse)) +
  geom_hline(aes(yintercept = 0))

# plot product access coefficient by year
by_year %>%
  ggplot(aes(x = year,
             y = prodaccessest)) +
  geom_point() +
  geom_pointrange(aes(ymax = prodaccessest + prodaccessse,
                      ymin = prodaccessest - prodaccessse)) +
  geom_hline(aes(yintercept = 0))

# plot clerk gender coefficient by year
by_year %>%
  ggplot(aes(x = year,
             y = clerkgenest)) +
  geom_point() +
  geom_pointrange(aes(ymax = clerkgenest + clerkgense,
                      ymin = clerkgenest - clerkgense)) +
  geom_hline(aes(yintercept = 0)) +
  labs(x = "Year",
       y = "Coefficient Estimate on Clerk Gender",
       title = "Figure 4: Effect of Clerk Gender Over Time")

# plot ask for ID coefficient by year
by_year %>%
  ggplot(aes(x = year,
             y = idest)) +
  geom_point() +
  geom_pointrange(aes(ymax = idest + idse,
                      ymin = idest - idse)) +
  geom_hline(aes(yintercept = 0)) +
  labs(x = "Year",
       y = "Coefficient Estimate on Asking for ID",
       title = "Figure 2: Effect of Asking for ID Over Time")
```

```r
# plot inspector age coefficient by year
by_year %>%
  ggplot(aes(x = year,
             y = ageest)) +
  geom_point() +
  geom_pointrange(aes(ymax = ageest + agese,
                      ymin = ageest - agese)) +
  geom_hline(aes(yintercept = 0))

# plot clerk age coefficient by year
by_year %>%
  ggplot(aes(x = year,
             y = clerkageest)) +
  geom_point() +
  geom_pointrange(aes(ymax = clerkageest + clerkagese,
                      ymin = clerkageest - clerkagese)) +
  geom_hline(aes(yintercept = 0))
```