Does the New York City Police Department Rely on Quotas?

Jonathan Auerbach
Department of Statistics
Columbia University

12/17/2017

Abstract

We investigate whether the New York City Police Department (NYPD) uses productivity targets or quotas to manage officers in contravention of New York State Law. The analysis is presented in three parts. First, we introduce the NYPD's employee evaluation system, and summarize the criticism that it constitutes a quota. Second, we describe a publically available dataset of traffic tickets issued by NYPD officers in 2014 and 2015. Finally, we propose a generative model to describe how officers write traffic tickets. The fitted model is consistent with the criticism that police officers substantially alter their ticket writing to coincide with departmental targets. We conclude by discussing the implications of these findings and offer directions for further research.

I. Introduction

Critics have periodically accused the New York City Police Department (NYPD) of relying on traffic ticket quotas to evaluate police officers and maintain productivity (Bacon 2009, pg. 97), (Rayman 2013, pg. 49, pg. 64), (Eterno and Silverman 2012, pg. 170). Such concerns are far from frivolous. Traffic ticket quotas are illegal in New York. Section 215-a of New York State Labor Law specifically prohibits NYPD supervisors from ordering an officer to write a predetermined number of tickets over a predetermined period of time.

Despite a lengthy history (Burrows and Wallace 1998), (Wallace 2017), accusations of quotas came to a head only recently, when the commander of the 75th precinct issued a memo directing officers to write ten tickets a month. The officers' union filed a grievance, and in 2006 an arbitrator ruled in their favor, ordering that "The city shall cease and desist from maintaining a vehicular ticket quota." (Fahim 2006) Yet instances of explicit ticket quotas continue to occur. The most notable came in 2010 when Officer Schoolcraft, a whistleblower from the 81st precinct, recorded his supervisor directing officers to write five traffic tickets a week each for drivers not wearing a seat belt, using a cellphone, double-parking and parking in a bus stop.

While the NYPD cannot legally set quotas, the law does not forbid the consideration of past productivity when assessing officers. In fact, the NYPD's evaluation system — Operations Order number 52, series 2011, titled "Quest for Excellence - Police Officer Performance Objectives" (hereafter abbreviated Q4E) — mandates it. Under Q4E, supervisors are directed to assess officers on the 7th, 14th and 21st of each month according to their productivity. At the end of the month, supervisors rate each officer as effective or ineffective. These supervisors are then rated by their own managers at the end of each quarter.

The pressure to do more with fewer resources leaves supervisors walking the fine line between incentivizing officer productivity and ordering it as Schoolcraft's supervisor did. The line is blurred when supervisors base evaluations on an officer's productivity relative to their peers. Spurring low performing officers to catch up to peers induces a "ticket race" between officers. Since supervisors control how officers are compared to the group, they control the pace of the race. The faster the race, the greater the number of tickets expected of each officer each period, effecting a de facto quota system (Dornsife 1978, pg. 3).

The legality of this practice has yet to be determined (Scheidlin 2013), but pressuring officers to increase productivity to compensate for past behavior reduces discretion and violates the spirit of the law as stated in the 1978 memorandum accompanying the original anti-quota bill:

> The police officer as well as the public, need not be put under the pressure of a mandatory ticket quota. Such a policy can only hurt the effectiveness of the police officer in the performance of his other duties while he must give the driving public a rash of summonses to meet the quota.

The ticket race is a plausible mechanism by which NYPD supervisors circumvent the anti-quota law and compel officers to give "a rash of summonses". However, the vast majority of evidence implicates only specific supervisors over short periods of time, without quantifying the magnitude or scope of the practice. Fortunately, traffic ticket data can be used determine if quota-like behavior exists across multiple precincts and time periods, or whether productivity targets reflect isolated incidents, confined to the times and places where physical evidence has been made available. The goal of the present analysis is to demonstrate how such an analysis might be performed.

II. Data

The primary dataset contains every recorded summons issued to drivers for a parking or moving violation in New York City between January 2014 and December 2015. It includes the id number of the issuing officer, their precinct, the date the ticket was written and the violation for which the ticket was written. The first five observations of the dataset are displayed in Table 1. Summary statistics for the first three columns are displayed in Table 2, and the most common traffic ticket violations are displayed in Table 3. Note that the ticket types from Officer Schoolcraft's recordings rank at 1, 4, 5 and 8 among the top ten ticket types.

```r
library("knitr")
library("timeDate")
library("scales")
library("tidyverse")
library("lubridate")
library("stringr")
library("rstan")
library("splines")

theme_set(theme_bw())
rstan_options(auto_write = TRUE)
options(mc.cores = parallel::detectCores())

download.file(str_c("https://raw.githubusercontent.com/jauerbach/",
                "Seventy-Seven-Precincts/master/data/tickets.csv.zip"),
              str_c(getwd(),"/tickets.csv.zip"))
tickets <- read_csv("tickets.csv.zip", col_types = cols(id = col_double()))

crashes <-
  read_csv(str_c("https://raw.githubusercontent.com/jauerbach/",
                 "Seventy-Seven-Precincts/master/data/crashes.csv"))
```

```r
tickets %>%
  head() %>%
  mutate(violation = substr(violation, 1, 48)) %>%
  kable(caption="Example Observations of Dataset")
```

Table 1: Example Observations of Dataset

id	command	date	violation
931369	40	2015-08-01	51-SIDEWALK
931369	40	2014-08-23	24-NO PARKING (EXC AUTH VEH)
931369	40	2014-08-23	45-TRAFFIC LANE
931369	40	2014-07-23	46-DOUBLE PARKING (COM PLATE)/46A-DOUBLE PARKING
951369	40	2015-07-23	09-BLOCKING THE BOX
931369	40	2015-07-24	66-DETACHED TRAILER

```r
tickets %>%
  select(-violation) %>%
  summary() %>%
  kable(caption="Summary of Dataset")
```

Table 2: Summary of Dataset

id	command	date
Min. : 0	Min. : 1.00	Min. :2014-01-01
1st Qu.:931632	1st Qu.: 28.00	1st Qu.:2014-06-19
Median :942676	Median : 66.00	Median :2014-12-11
Mean :934616	Mean : 61.94	Mean :2014-12-21
3rd Qu.:949969	3rd Qu.:100.00	3rd Qu.:2015-06-22
Max. :999992	Max. :123.00	Max. :2015-12-31

```r
tickets %>%
  count(violation) %>%
  arrange(desc(n)) %>%
  mutate(violation = substr(violation, 1, 48)) %>%
  head(n = 10) %>%
  kable(caption= "Most Common Traffic Tickets")
```

Table 3: Most Common Traffic Tickets

violation	n
46-DOUBLE PARKING (COM PLATE)/46A-DOUBLE PARKING	457397
1110A-DISOBEYED TRAFFIC DEVICE	234775
40-FIRE HYDRANT	166296
19-NO STAND (BUS STOP)	144019
1225C2A-OPERATING MV MOBILE PHONE	138590
14-NO STANDING	133958
20-NO PARKING (COM PLATE)/20A-NO PARKING (NON-CO	123725
1229C3A-NO LAP/SHOULDER HARNESS OR DJ VIO	100726
37512A2-SIDEWINGS/SIDEWINDOWS/NON/TRANSPARENT	92913
98-OBSTRUCTING DRIVEWAY	84219

The unit of this analysis is the number of tickets written by each officer each day, and the following code block aggregates observations by officer and day. Days of the month are grouped into four review periods depending on their proximity to the Q4E evaluations on the 7th, 14th, 21st and end of the month. For example, the first review period contains the first seven days of the month, the second review period contains the eighth of the month to the fourteenth, and so on.

Some officers are only represented for a portion of the two-year sample period. We assume an officer is eligible to write tickets for every day in a month if at least one ticket was written by the officer that month. Otherwise, the officer is excluded for the entire month. The only exception is days during which no ticket is written in the entire precinct. Those days are excluded from the analysis.

```r
# calculate the number of tickets per day.
by_day <- tickets %>%
  mutate(month_year = format(date, "%m-%Y")) %>%
  group_by(id, command, date, month_year) %>%
  summarise(daily_tickets = n()) %>%
  ungroup()

# add zeros and period mean and max, assuming at least one ticket written in
## command during month for a day to be eligible for writing tickets
by_day <- by_day %>%
  group_by(command, month_year) %>%
  nest() %>%
  mutate(elig_days = map(data, function(df) df %>% expand(id, date)),
         data_aug = map2(elig_days, data, left_join, by = c("id","date"))) %>%
  select(command, month_year, data_aug) %>%
  unnest() %>%
  mutate(daily_tickets = ifelse(is.na(daily_tickets), 0, daily_tickets))

# add mean tickets and max tickets for each command each review period. review
## periods are on the 7th, 14th, 21st and end of each month
by_day <- by_day %>%
```

```r
  mutate(period = cut(parse_number(format(date, "%d")),
                      breaks = c(1,7,14,21,31), labels = FALSE,
                      right = TRUE, include.lowest = TRUE)) %>%
  group_by(command, month_year, period) %>%
  nest() %>%
  mutate(mean = map_dbl(data, function(df) df %>%
                          group_by(id) %>%
                          summarise(sum(daily_tickets)) %>%
                          pull %>%
                          mean()),
         median = map_dbl(data, function(df) df %>%
                          group_by(id) %>%
                          summarise(sum(daily_tickets)) %>%
                          pull %>%
                          median()),
         max = map_dbl(data, function(df) df %>%
                         group_by(id) %>%
                         summarise(sum(daily_tickets)) %>%
                         pull %>%
                         max())) %>%
  unnest()

#add number of tickets per officer per period
by_day <- by_day %>%
  left_join(by_day %>%
              group_by(command, month_year, period, id) %>%
              summarise(period_tickets = sum(daily_tickets)) %>%
              ungroup(),
            by = c("command", "month_year", "period", "id"))

#add number of tickets, command max/mean in pervious period
by_day <- by_day %>%
  left_join(by_day %>%
              select(-daily_tickets, - date) %>%
              unique() %>%
              mutate(period = period + 1) %>%
              rename(mean_prev = mean,
                     median_prev = median,
                     max_prev = max,
                     tickets_prev = period_tickets),
            by = c("command", "month_year", "period", "id")) %>%
  mutate(mean_prev = ifelse(is.na(mean_prev), 0, mean_prev),
         median_prev = ifelse(is.na(median_prev), 0, mean_prev),
         max_prev = ifelse(is.na(max_prev), 0, max_prev),
         tickets_prev = ifelse(is.na(tickets_prev), 0, tickets_prev))

#add US Holidays
by_day <- by_day %>%
  left_join(sapply(c("USNewYearsDay", "USMLKingsBirthday", "USWashingtonsBirthday",
                     "USMemorialDay", "USIndependenceDay", "USLaborDay",
                     "USColumbusDay", "USElectionDay", "USVeteransDay",
                     "USThanksgivingDay", "USChristmasDay"),
                   function(x) as.Date(holiday(2014:2015, x))) %>%
```

```r
              as_tibble() %>%
              gather(key = "holidays", value = "date") %>%
              mutate(date = as.Date(date, origin = "1970-01-01")),
            by = "date") %>%
  mutate(holidays = ifelse(is.na(holidays), "USNone", holidays))
```

III. Preliminary Evidence

For computational reasons, we limit the present analysis to Schoolcraft's Precinct, Precinct 81.

The data indicate that more tickets are written in the second half of the month, and this increase comes from officers that are behind their peers. Figure 1.1 demonstrates that the average number of tickets written per day increases from the first half to the second half of the month. Here, the second half of the month begins on the 15th, following the second evaluation date under Q4E. The increase between the two halves is six percent.

```r
by_day %>%
  filter(command == 81) %>%
  mutate(half = ifelse(period > 2,
                       "Second Half \nof Month",
                       "First Half \nof Month")) %>%
  group_by(half) %>%
  summarise(y = sum(daily_tickets)) %>%
  mutate(y = ifelse(half == "First Half \nof Month",
                    y / (24 * 14),
                    y * 12 / (24 * 198))) %>%
  ggplot() +
    aes(half, weight = y) +
    geom_bar() +
    theme(legend.position = "bottom") +
    scale_y_continuous(labels = comma) +
    labs(y = "average number of tickets written in 2014-15", x = "",
         fill = "",
         title="Figure 1.1:
         the ticketing rate increases six percent over the month ...")
```

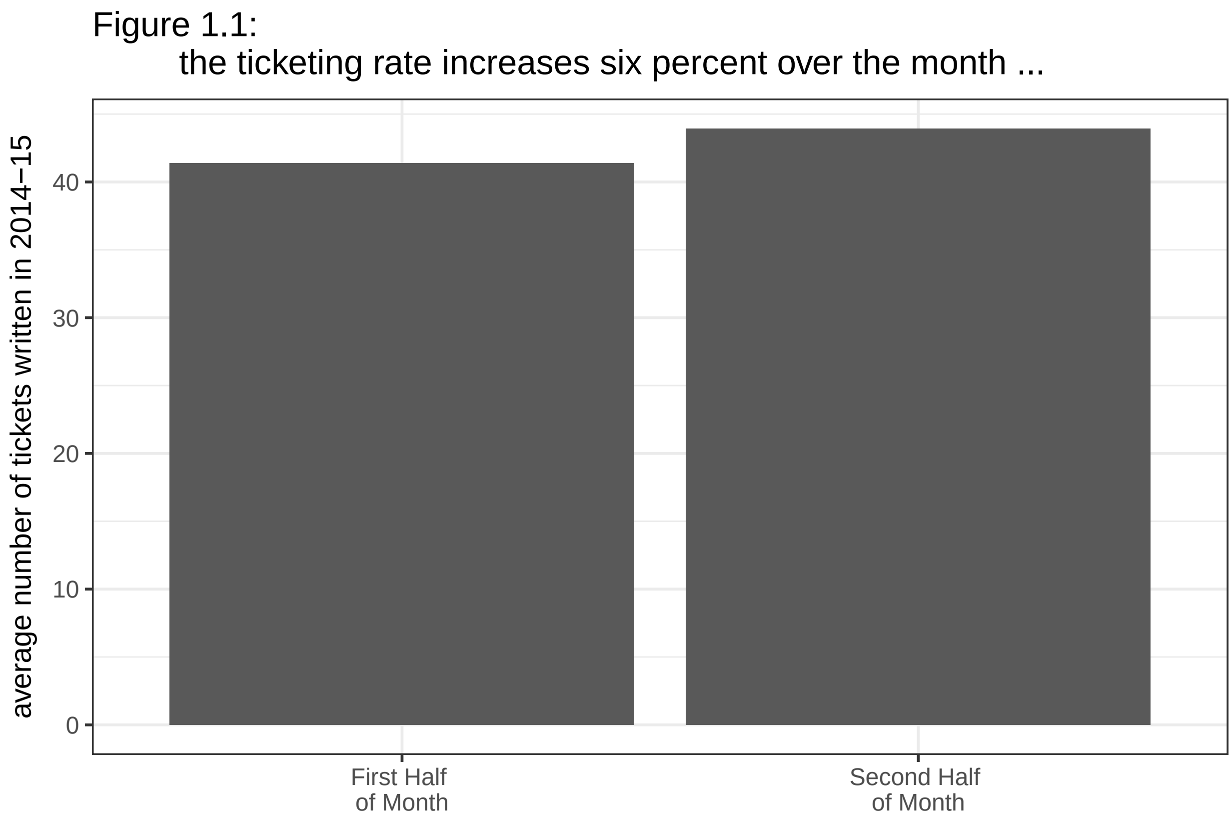

Figure 1.2 demonstrates that the increase in Figure 1.1 is unlikely the result of driver behavior since the average number of vehicle crashes does not increase over this period. If changes in exposure were causing the increase in Figure 1.1, more crashes would be expected as well.

```
crashes %>%
  filter(Precinct == 81) %>%
  mutate(period = cut(parse_number(format(DATE, "%d")),
                      breaks = c(1,7,14,21,31), labels = FALSE,
                      right = TRUE, include.lowest = TRUE),
         half = ifelse(period > 2,
                       "Second Half \nof Month",
                       "First Half \nof Month")) %>%
  group_by(half) %>%
  summarise(y = sum(n)) %>%
  mutate(y = ifelse(half == "First Half \nof Month",
                    y / (24 * 14),
                    y * 12 / (24 * 198))) %>%
  ggplot() +
    aes(half, weight = y) +
    geom_bar() +
    theme(legend.position = "bottom") +
    scale_y_continuous(labels = comma) +
    labs(y = "average number of collisions in 2014-15", x = "",
         fill = "",
         title="Figure 1.2:
         ...while the number of collisions remains unchanged")
```

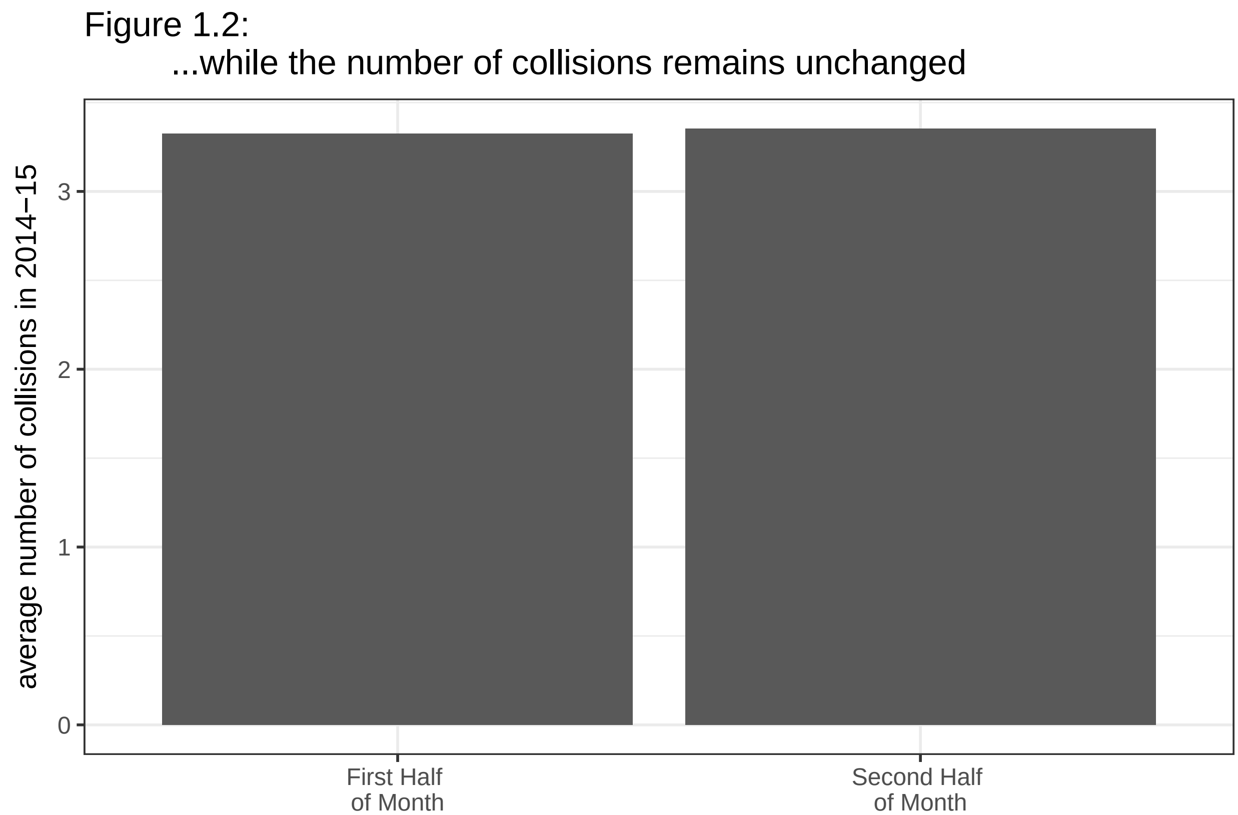

Figure 1.3 shows that the increase in the number of tickets issued is due entirely to officers who have below median productivity. Since there are more days after the 15th than before it, a constant number of traffic ticket translates to a decrease in the average number of tickets written. Much of that decrease occurs in the fourth period as displayed in Figure 1.4.

```
by_day %>%
  filter(command == 81) %>%
  mutate(prod = ifelse(tickets_prev >= median_prev,
                       "Officer >= Median","Officer < Median"),
         half = ifelse(period > 2,
                       "Second Half \nof Month",
                       "First Half \nof Month")) %>%
  group_by(half, prod) %>%
  summarise(y = sum(daily_tickets)) %>%
  ggplot() +
    aes(half, weight = y, fill = factor(prod)) +
    geom_bar() +
    theme(legend.position = "bottom") +
    scale_y_continuous(labels = comma) +
    labs(y = "all traffic tickets written in 2014-15", x = "period", fill = "",
         title="Figure 1.3:
         officers behind their peers account for the entire increase in the
         number of tickets written in the second half of the month ...")
```

Figure 1.3:
officers behind their peers account for the entire increase in the number of tickets written in the second half of the month ...

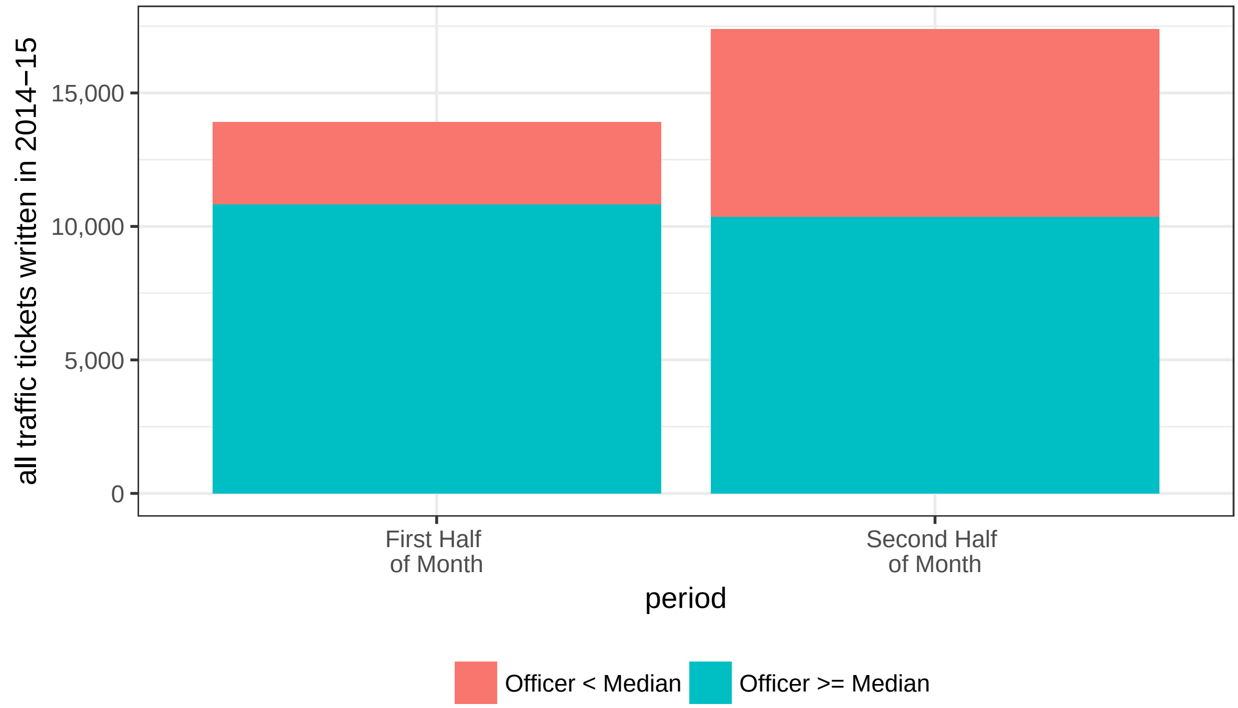


```r
by_day %>%
  filter(command == 81,
         period != 1) %>%
  mutate(prod = ifelse(tickets_prev > median_prev,
                       "Officer >= Median","Officer < Median")) %>%
  group_by(period, prod) %>%
  summarise(y = sum(daily_tickets)) %>%
  mutate(y = ifelse(period == 4,
                    y * 12 / (24 * 134),
                    y / (24 * 7))) %>%
  ggplot() +
    aes(period, weight = y, fill = factor(prod)) +
    geom_bar(position = "dodge") +
    theme(legend.position = "bottom") +
    scale_y_continuous(labels = comma) +
    labs(y = "average number of tickets written in 2014-15",
         x = "period", fill = "",
         title="Figure 1.4:
         ... but it is officers ahead of their peers that are drastically
         reducing their ticket writing rate")
```

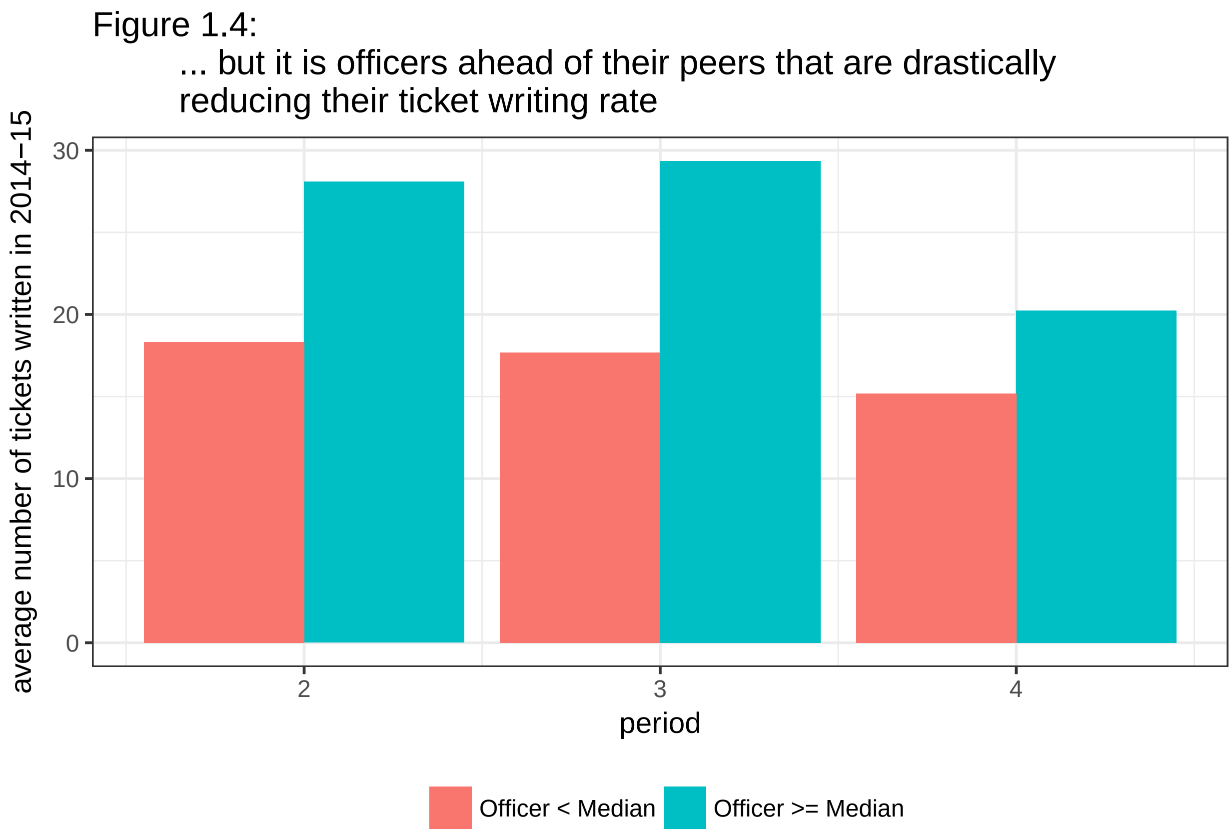

The data are consistent with the criticism that officers winning the ticket race reduce the number of tickets written each period and officers losing the ticket race increase the number of tickets (Dornsife 1978, pg. 3). The problem with the analysis so far is that these changes could be a selection effect. If officer activity were random each period, one would expect an increase after selecting the lowest performing officers and a decrease after selecting the highest performing officers. Productivity should regress towards the mean in either case by default, as an artifact of the selection process.

A model is needed to fully evaluate the working theory. However, any model is complicated by the fact that the majority of officers wrote a minority of traffic tickets. As can be seen in the Figure 1.5, over the two year study period, ninety percent of officers wrote one or fewer tickets each day, while the top one percent of officers wrote more than one thousand tickets. However, these officers are responsible for the increase in the overall ticket rate and their behavior, not the behavior of the most common ticket writers, is of primary interest.

In this sense, the data is sparse, and information is spread out among a large number of relatively rare events. Determining whether officers have increased their ticketing rate fits well within the Bayesian hierarchical model paradigm for this reason: information can be pooled across officers to make generalizing statements of officer behavior.

```
by_day %>%
  filter(command == 81) %>%
  ggplot() +
    aes(daily_tickets) +
    geom_histogram() +
    scale_y_log10(labels = comma) +
    labs(x = "number of tickets written per officer per day in 2014-15",
         y = expression(log[10] ~ " number of officers"),
         title="Figure 1.5:
         the rate officers write tickets has a heavy tail")
```

Figure 1.5:
the rate officers write tickets has a heavy tail

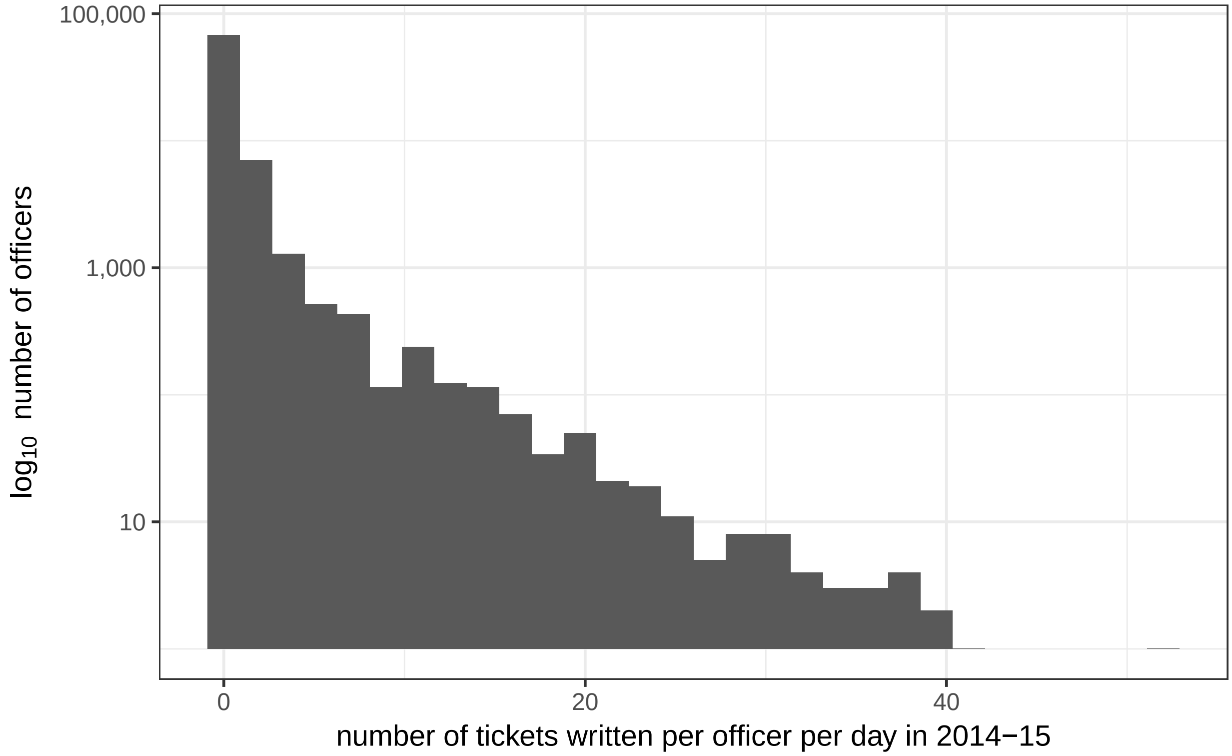

IV. Model

We begin by combining two log-linear models for contingency tables. Let y_{od} be the number of traffic tickets written by the o^{th} officer on the d^{th} day, and let z_d be the number of crashes in the officers' precinct on the d^{th} day. We include an effect for each day of the week, w, holiday, h, period, p, and month m. Week, holiday, period and month effects reflect various administrative policies. These effects are then pooled as follows:

$$
\begin{aligned}
\sigma. &\sim \text{ Chi-Squared}(1) \\
\beta. &\sim \text{Normal}(0, \sigma.) \\
\epsilon_d &\sim \text{Normal}(\beta_0, 1) \\
y_{od} &\sim \text{Poisson}(\exp\{\beta_o + \beta_w + \beta_p + \beta_h + \beta_m + \sigma_\epsilon \epsilon_d\}) \\
z_d &\sim \text{Poisson}(\exp\{\beta_0 + \sigma_0 \epsilon_d\})
\end{aligned}
$$

The choice of Chi-Squared(1) priors was arbitrary. The model was also fit with Normal$(0, 1)$ priors on the standard deviation parameters, and little change was observed in the outcomes.

```
model_data <- by_day %>%
   filter(command ==  81) %>%
  mutate(id2 = as.numeric(as.factor(id)),
         date2 = as.numeric(as.factor(date)),
         months = months(date),
         months2 = as.numeric(format(date, "%m")),
```

```
         weekdays = weekdays(date),
         weekdays2 = wday(date),
         holidays2 = as.numeric(as.factor(holidays))) %>%
  with(list(N = nrow(by_day %>% filter(command ==  81)),
            O = length(unique(id2)),
            W = length(unique(weekdays2)),
            M = length(unique(months2)),
            H = length(unique(holidays2)),
            P = length(unique(period)),
            officer = id2,
            week = weekdays2,
            month= as.numeric(months2),
            holiday= as.numeric(holidays2),
            period = as.numeric(period),
            day = date2,
            y = daily_tickets,
            D = crashes %>%
                filter(Precinct == 81) %>%
                semi_join(by_day %>% filter(command ==  81),
                          by = c("DATE" = "date")) %>%
                  select(n) %>%
                  pull() %>%
                  length() ,
            z = crashes %>%
                filter(Precinct == 81) %>%
                semi_join(by_day %>% filter(command ==  81),
                          by = c("DATE" = "date")) %>%
                  select(n) %>%
                  pull(),
            X = ifelse(tickets_prev - median_prev > 10, 10,
                       tickets_prev - median_prev)))

fit1 <- sampling(stan_model(file = "model1.stan"),
                 data = model_data,
                 iter = 500,
                 control = list(max_treedepth = 15, adapt_delta = 0.9))
fit1_summary <- summary(fit1)
```

We use RStan (Stan Development Team 2016) to run four chains for 500 iterations each that approximate samples from the posterior distribution of the above model. The first 250 iterations are discarded in warm-up, leaving one thousand posterior draws to estimate the parameters. After increasing the max_treedepth and adapt_delta arguments, no notable issues occurred during sample. An overview of parameter estimates is provided in Figure 2, in which the average ticket writing of officers in Precinct 81 is decomposed into baseline, week, period and month effects (holiday effects not shown). Period and month patterns are evident, showing an increase before evaluations at the end of the month and quarter respectively. We also observe the decline around Christmas of 2014, which was the result of an alleged slowdown.

```
var <- rownames(fit1_summary[[1]]); est <- fit1_summary[[1]][,1]
dfs <- list(
    date = by_day %>%
            filter(command ==  81) %>%
            transmute(date,
                      per = period,
                      week_day = weekdays(date),
```

```r
                      mnth = month(date),
                      daily_tickets) %>%
                     group_by(date, per, week_day, mnth) %>%
             summarise(`log tickets` = log(mean(daily_tickets))),
    base = crashes %>%
              filter(Precinct == 81) %>%
              semi_join(by_day %>%
              filter(command ==  81), by = c("DATE" = "date")) %>%
              transmute(date = DATE,
                        baseline = est[substr(var,1,7) == "sigma_0"] *
                                      est[substr(var,1,6) == "epsilo"]
                        ),
    day = tibble(`week day` = est[substr(var,1,7) == "sigma_2"] *
                     est[substr(var,1,6) == "beta_2"],
                 week_day = c("Sunday", "Monday", "Tuesday", "Wednesday",
                              "Thursday", "Friday", "Saturday")),
    period = tibble(period = est[substr(var,1,7) == "sigma_5"] *
                      est[substr(var,1,6) == "beta_5"], per = 1:4),
    month = tibble(month = est[substr(var,1,7) == "sigma_3"] *
                     est[substr(var,1,6) == "beta_3"], mnth = 1:12))

reduce(dfs, left_join) %>%
    gather("effect", "value", `log tickets`, baseline, `week day`, period, month) %>%
    ggplot() +
        aes(date, value, color = effect) +
        geom_line() +
        facet_grid(factor(effect,
          levels = c("log tickets", "baseline", "week day", "period", "month")) ~ .,
          scales = "free") +
        theme_bw() +
        theme(legend.position = "none") +
        scale_x_date(date_labels = "%b",
              breaks = parse_date(c("2014-01-01","2014-04-01","2014-07-01",
                                    "2014-10-01", "2015-01-01","2015-04-01",
                                     "2015-07-01","2015-10-01", "2016-01-01"))) +
             labs(x = "", y = "",
         title="Figure 2:
         decomposition of log ticket rate implied by model")
```

Figure 2:
decomposition of log ticket rate implied by model

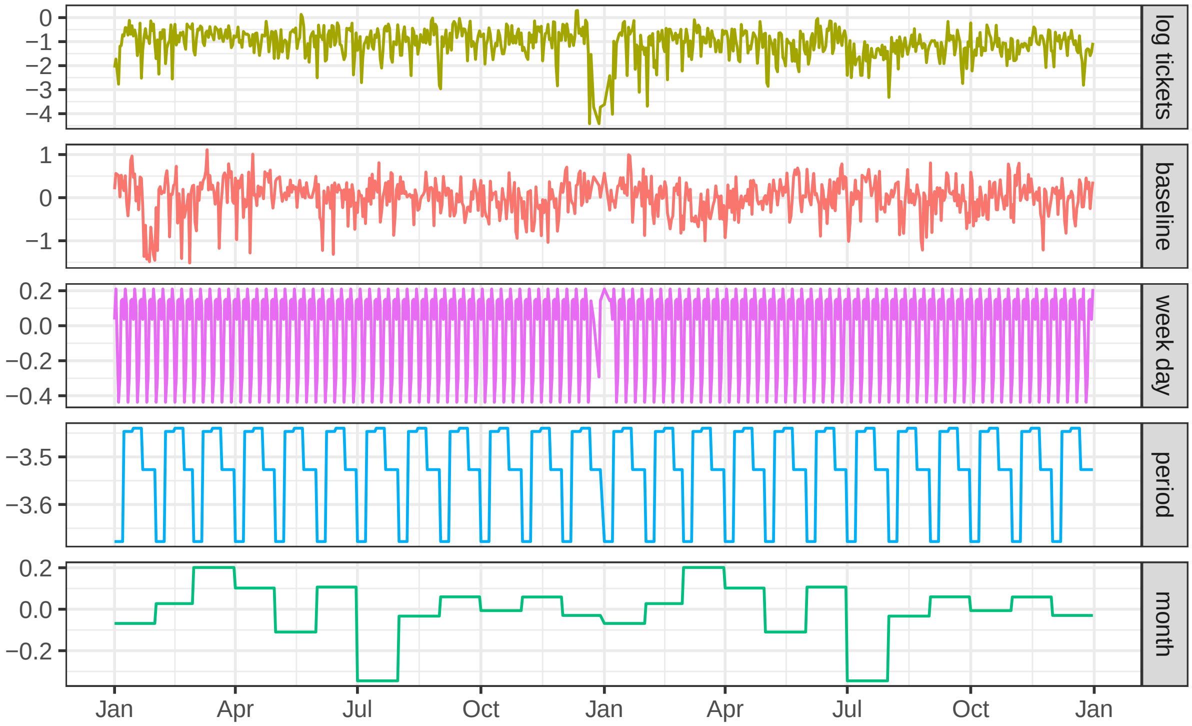

Uncertainty intervals for the effects (Gelman 2005), (Gelman and Hill 2006) are depicted in Figures 3.2 through 3.4. Fat black lines correspond to 50 percent uncertainty intervals, while thin black lines correspond to 95 percent intervals. We interpret these intervals as depicting plausible and possible values of the parameters.

Weekends and Wednesdays have lower ticket writing rates even after adjusting for the number of crashes. This likely corresponds with alternate side parking rules that disproportionately effect Monday/Wednesday and Tuesday/Fridays. March, June and September have higher ticket writing rates, which makes sense because, at the end of these months, supervisors are themselves evaluated. July and December have lower ticket writing rates. These are the traditional vacation months. November may be higher because of a pre-vacation push before the end of the year.

Periods are not statistically different from each other after adjusting for the other covariates. This is not necessarily inconsistent with quota-like behavior since we would expect officers ahead in the ticket race to decrease their activity while officers behind would increase it. This reflection motivates the following refinement to the model.

```
plot_coef <- function(summary, coef, label = NULL, str = 6) {
  if(is.null(label)) {
  var_name = rownames(summary$summary)[substr(rownames(summary$summary),1,str) == coef]
  } else {
    var_name = label
  }
  data.frame(
    name = var_name,
    mean = summary$summary[substr(rownames(summary$summary),1,str) == coef, "mean"],
    lower1 = summary$summary[substr(rownames(summary$summary),1,str) == coef, "mean"] -
      summary$summary[substr(rownames(summary$summary),1,str) == coef, "sd"],
    upper1 = summary$summary[substr(rownames(summary$summary),1,str) == coef, "mean"] +
      summary$summary[substr(rownames(summary$summary),1,str) == coef, "sd"],
```

```
    lower2 = summary$summary[substr(rownames(summary$summary),1,str) == coef, "mean"] -
      2 * summary$summary[substr(rownames(summary$summary),1,str) == coef, "sd"],
    upper2 = summary$summary[substr(rownames(summary$summary),1,str) == coef, "mean"] +
      2 * summary$summary[substr(rownames(summary$summary),1,str) == coef, "sd"]) %>%
    ggplot() +
      aes(x = name, y = exp(mean), ymin = exp(lower1), ymax = exp(upper1)) +
    geom_linerange(size = 2) +
    geom_linerange(aes(ymin = exp(lower2), ymax = exp(upper2))) +
    coord_flip()
}

qplot(fit1_summary$summary[
  substr(rownames(fit1_summary$summary),1,6) == "beta_1", "mean"]) +
  labs(x = expression(e^beta[officer]), y = "",
         title="Figure 3.1:
         distribution of officer effects in Precinct 81")
```

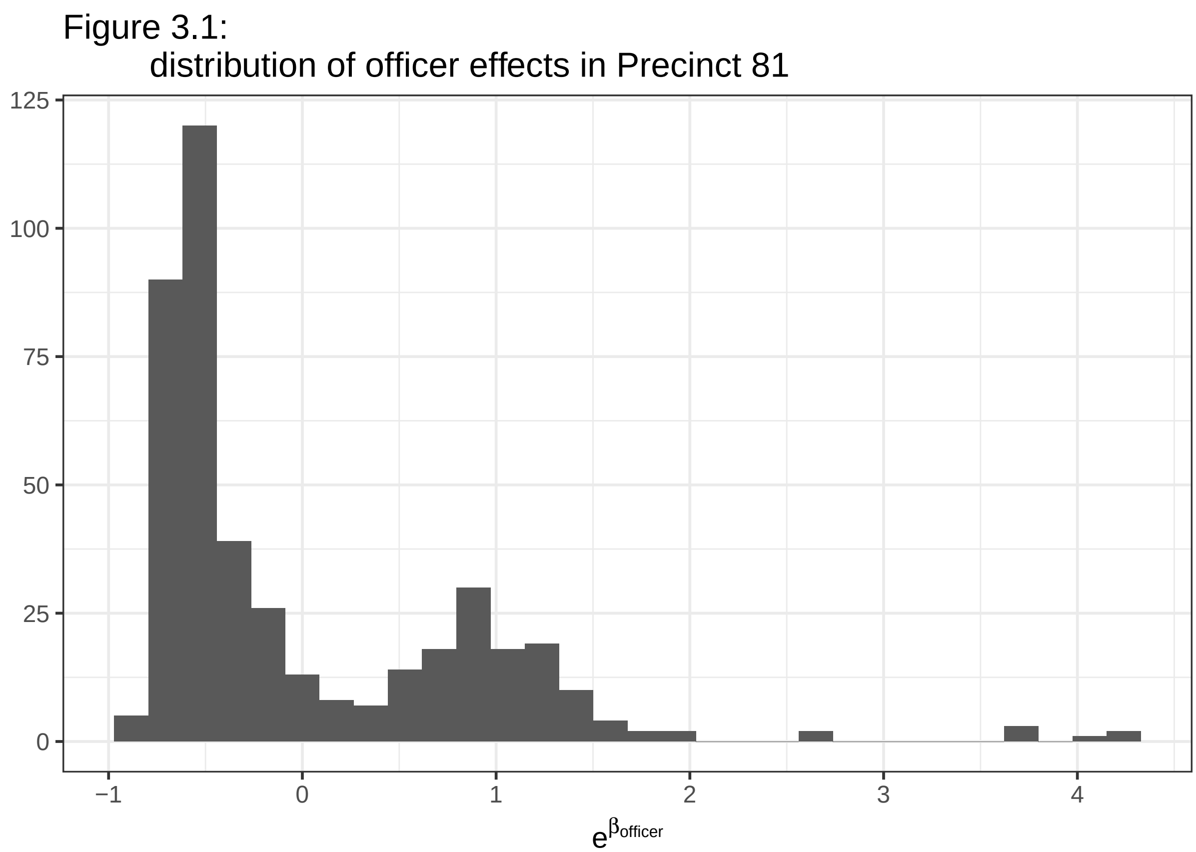


```
plot_coef(summary = fit1_summary, coef = "beta_2",
          label = factor(c("Sunday", "Monday", "Tuesday", "Wednesday",
                           "Thursday", "Friday", "Saturday"),
                         levels = weekdays(x=as.Date(seq(7),
                                                     origin="1950-01-01")))) +
  labs(x ="", y = expression(e^beta[w]),
         title="Figure 3.2:
         day of week effects in Precinct 81")
```

Figure 3.2:
day of week effects in Precinct 81

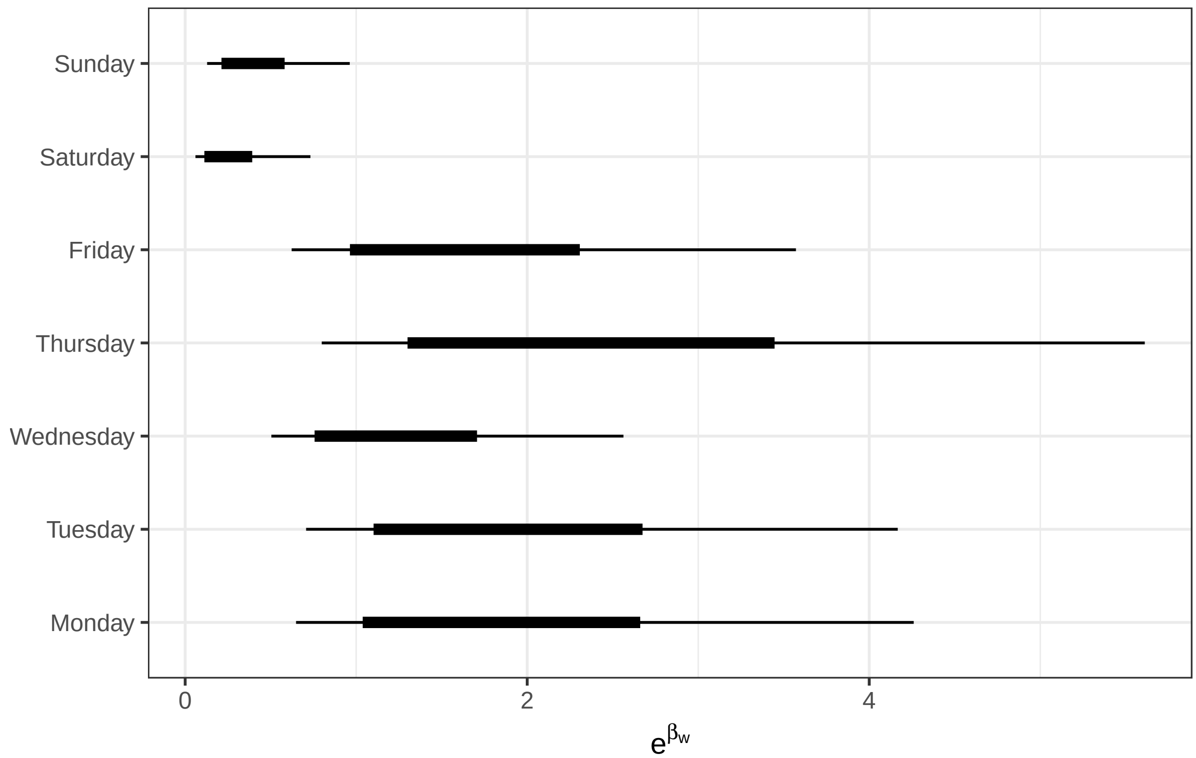


```r
plot_coef(summary = fit1_summary, coef = "beta_5", label = 1:4) +
  labs(x ="period effect", y = expression(e^beta[p]),
       title="Figure 3.3:
         period effects in Precinct 81")
```

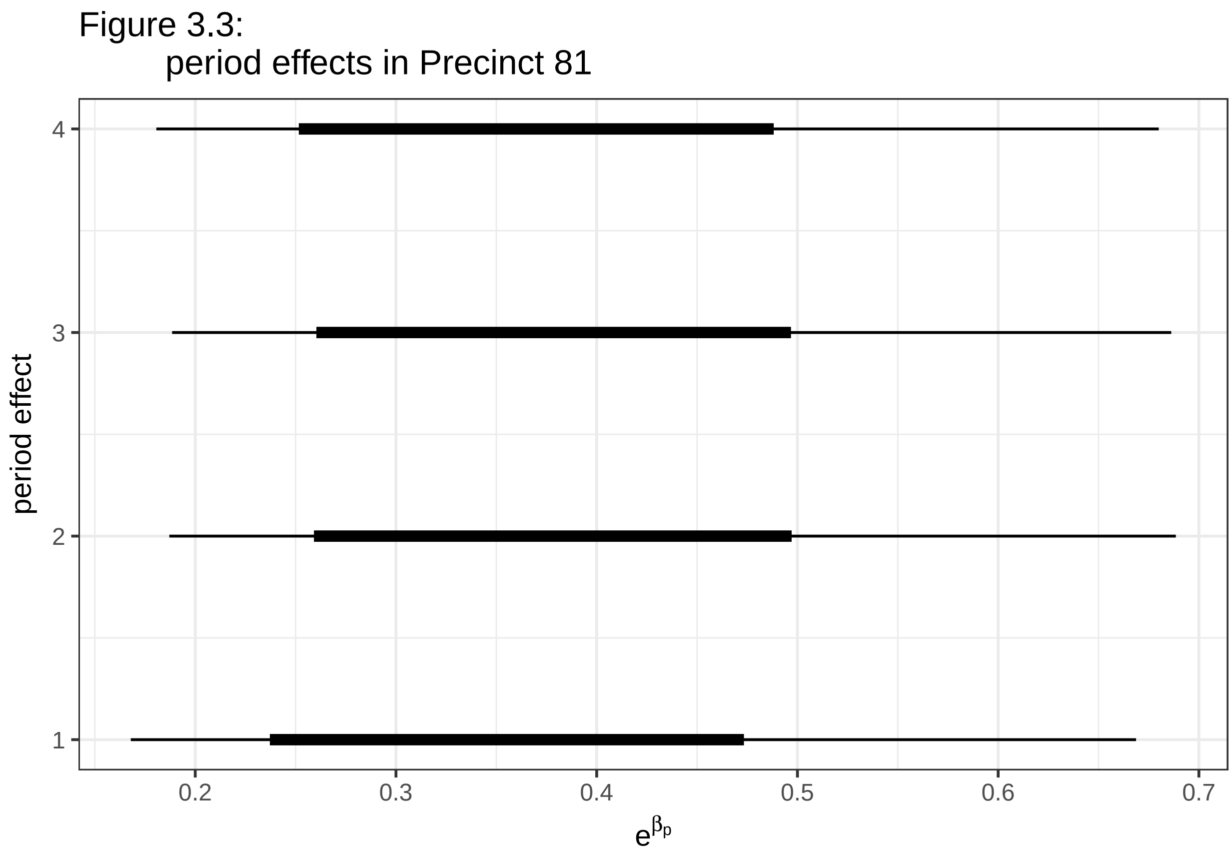


```
plot_coef(summary = fit1_summary, coef = "beta_4",
          label =   factor(c("Columbus Day", "Election Day",
                             "Independence Day", "Labor Day",
                             "Memorial Day", " M L Kings Birthday",
                             "New Years Day", "None", "Thanksgiving Day",
                             "Veterans Day", "Washingtons Birthday")) %>%
    (function(x) factor(x, levels = x[c(7, 6, 11, 5, 3, 4, 1, 10, 2, 9, 8)]))) +
  labs(x ="", y = expression(e^beta[h]),
  title="Figure 3.4:
        holiday effects in Precinct 81")
```

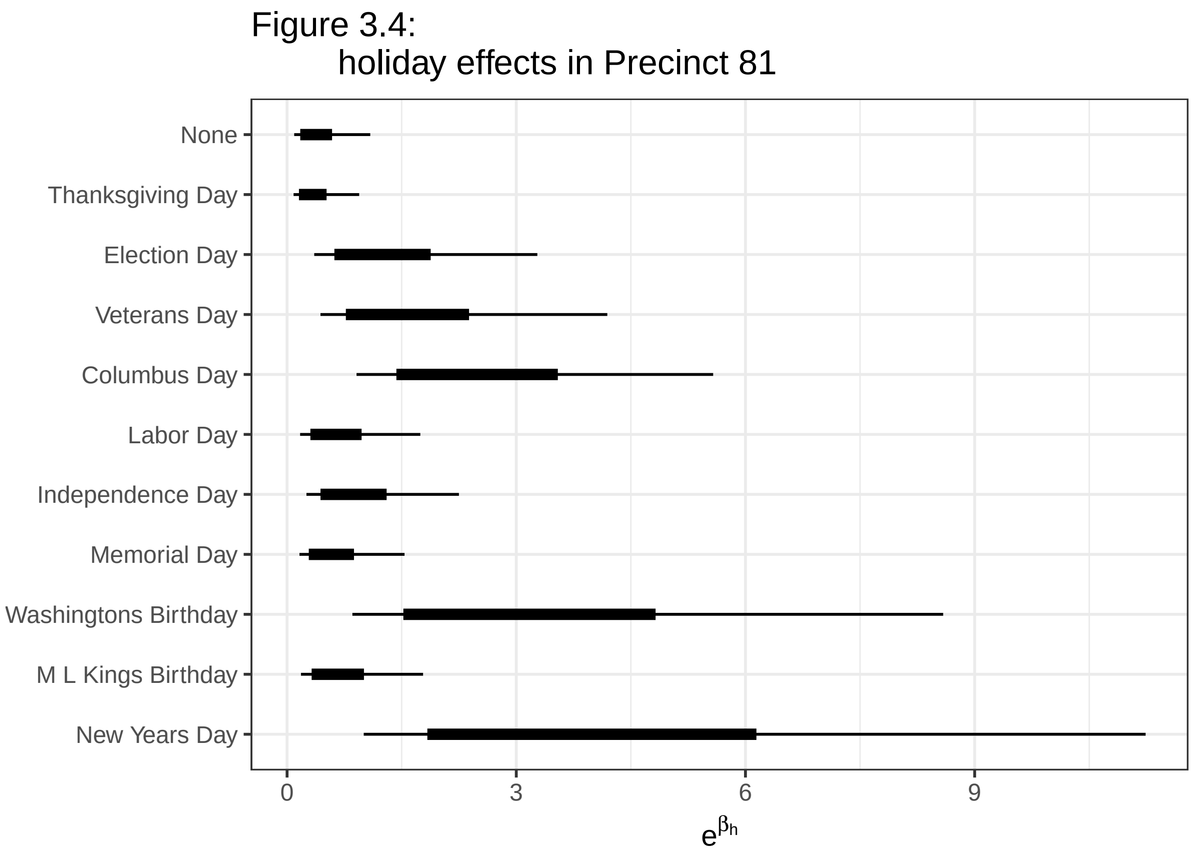


```
plot_coef(summary = fit1_summary, coef = "beta_3",
          label = month(x=as.Date(29 * seq(12), origin="1950-01-01"),
                        label = TRUE)) +
  labs(x ="", y = expression(e^beta[m]),
        title="Figure 3.5:
         month effects in Precinct 81")
```

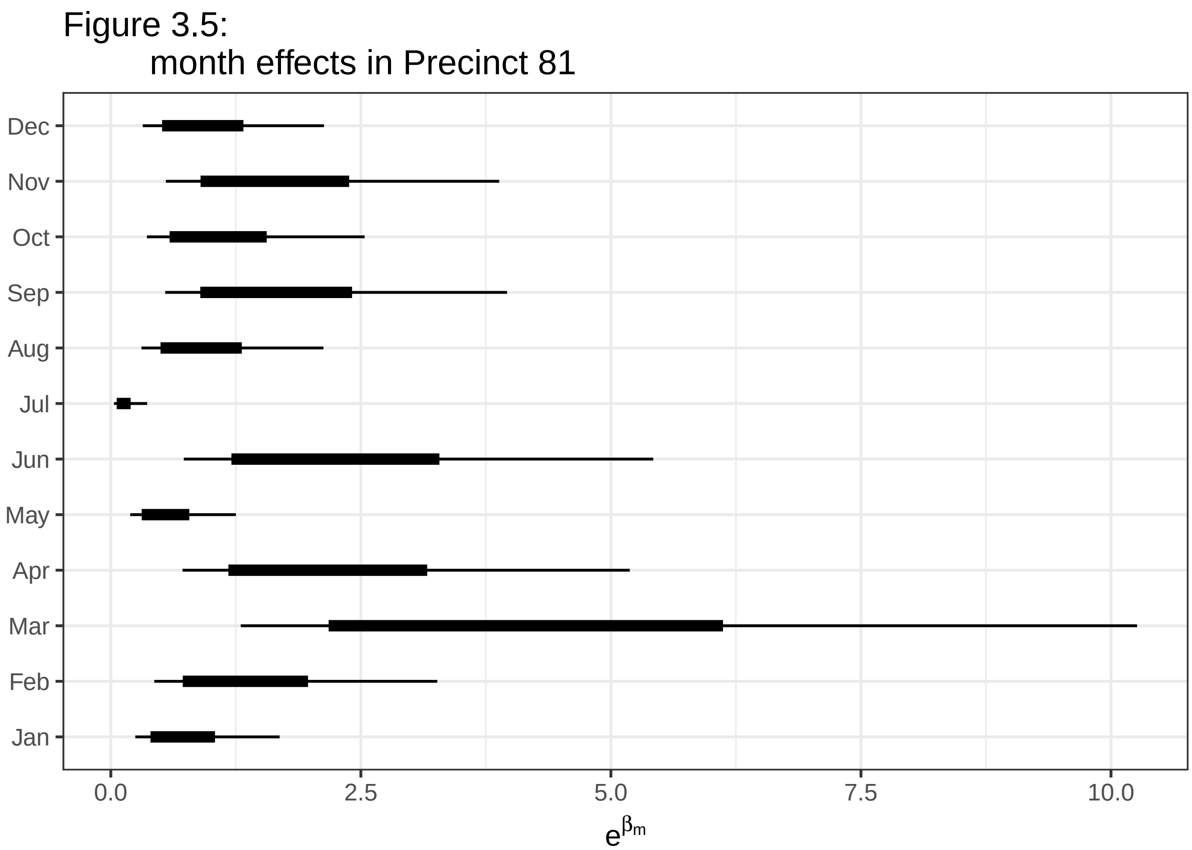

We now replace the period effect, β_p, with $f(rp_{o,p-1})$ where $rp_{o,p-1}$ is the relative productivity of officer o in the previous period $p-1$. Relative productivity is defined to be the number of tickets written by the officer in the previous period minus the median number of tickets written by all officers of that period in the same precinct. When $p = 1$, rp is taken to be zero for all officers. f is assumed to be smooth and approximated by a penalized B-spline, with code appropriated, with minor changes, from Kharratzadeh (2017). For identifiability reasons, officers more than 10 tickets ahead are considered 10 tickets ahead.

$$
\begin{aligned}
\sigma_. &\sim \text{ Chi-Squared}(1) \\
\beta_. &\sim \text{Normal}(0, \sigma_.) \\
\epsilon_d &\sim \text{Normal}(\beta_0, 1) \\
y_{od} &\sim \text{Poisson}(\exp\{\beta_o + \beta_w + f(rp_{o,p-1}) + \beta_h + \beta_m + \sigma_\epsilon \epsilon_d\}) \\
z_d &\sim \text{Poisson}(\exp\{\beta_0 + \sigma_0 \epsilon_d\})
\end{aligned}
$$

A plot of f is displayed below with standard error lines. The shape of f is quite dramatic, suggesting that past ticket writing behavior drives future behavior. If an officer is five tickets below the median, their ticketing rate is expected to be ten percent above the average in the following period. If an officer is at the median, their ticketing rate is expected to be twenty percent below the average. Finally, if the officer is five tickets above the median, the ticketing rate is thirty percent below the average.

```
knots <- c(-5, -3, -1, 0, 1, 3, 5, 7, 10)
num_knots <- length(knots)
spline_degree <- 3

fit2 <- sampling(stan_model(file = "splines.stan"), data = model_data,
```

```r
                 chains = 4, iter = 500,
                 control = list(max_treedepth = 15, adapt_delta = 0.9))
fit2_summary <- summary(fit2)
```

```r
f_coef <- fit2_summary$summary[substr(rownames(fit2_summary$summary),1,6) == "lambda", 1]
f_sd <- fit2_summary$summary[substr(rownames(fit2_summary$summary),1,6) == "lambda", 3]

qplot(model_data$X, exp(f_coef), geom= "line") +
  geom_line(aes(x, exp(y)),
            data = data.frame(x = model_data$X, y = f_coef + f_sd),
            linetype = 2) +
  geom_line(aes(x, exp(y)),
            data = data.frame(x = model_data$X, y = f_coef - f_sd),
            linetype = 2) +
  xlim(-4.5, 4.5) +
  labs(x = expression(rp["o,p-1"]), y = expression(e^f(rp["o,p-1"])),
       title="Figure 4:
         effect of position in ticket race in Precinct 81")
```

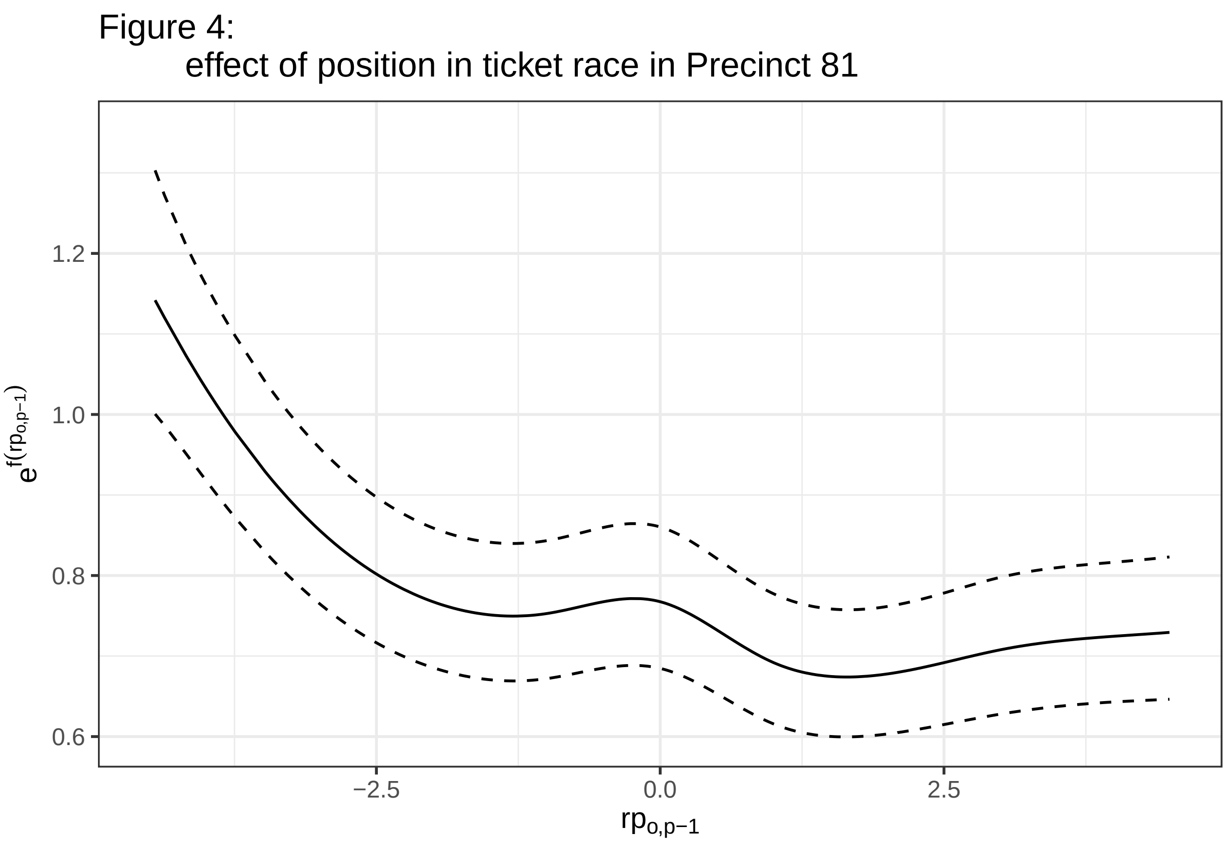

V. Conclusion

The fitted model is consistent with the criticism that police officers alter their ticket writing to coincide with departmental targets. Figure 4 indicates that the behavior of Precinct 81 officers in past periods influences their future behavior, and that this influence is quite dramatic, accounting for as much as a fifty percentage point spread in the rate officers write tickets. Were the ticketing rate solely a function of road conditions, we would expect past behavior to have no effect on future behavior, and little change should be observed in

Figures 1 or 4.

We conclude by making several, distinct remarks. We first point out that other theories could explain the pattern observed in the data. For example, the relationship depicted in Figure 4 could be a consequence of how the NYPD deploys officers. Rotating schedules each period would result in low ticketing officers one period increasing their productivity in the next. This seems unlikely, however, since nothing in Q4E or the literature suggests that scheduling takes place in a manner that regularly coincides with Q4E review dates. Nevertheless, we could account for this explanation by expanding the model to include interactions between officer and day of week or period effects. In addition, pooling across multiple precincts would also provide evidence that the estimated relationship is not due to precinct specific practices.

Regardless, the increases observed in Figures 1 and 4 are evidence of increased officer activity. We also point out that this increase is barely perceptible to the average driver. Even the largest estimated increase, a twenty-five percent change from officers five tickets below their peers, would be difficult for a driver to perceive. One would need to receive more than 250 tickets from an officer to reject the hypothesis that their ticket writing rate is not constant between the two periods.

The twenty-five percent increase is significant because it represents a systemic shift in activity, and at issue is the quality of that activity. A quota system is illegal because the micromanagement of police officers reduces discretion. A de facto quota system, based on a "ticket race", serves the same purpose. History has revealed that the indiscriminate enforcement of minor violations can provoke confrontations between police and the community and erode the public trust necessary for combating serious offenses, which is in direct opposition to the officer's mission. In extreme cases, a ticket can ruin entire families by triggering the parole violation, incarceration or deportation of an otherwise upstanding resident.

Yet performance goals need not micromanage officers. In the landmark opinion that rejected the implementation of NYPD's Stop and Frisk procedure, the judge found the use of performance goals under Q4E "created pressure to carry out stops, without any system for monitoring the constitutionality of those stops. However, the use of performance goals in relation to stops may be appropriate, once an effective system for ensuring constitutionality is in place." (Scheidlin 2013, pg. 17) We believe data analysis will be an important part of any system, and we see this work as contributing towards a conversation of how such an analysis could be performed.

VI. References

Bacon, Paul. 2009. *Bad Cop: New York's Least Likely Police Officer Tells All.* Bloomsbury Publishing USA.

Burrows, Edwin G, and Mike Wallace. 1998. *Gotham: A History of New York City to 1898.* Oxford University Press.

Dornsife, Rod. 1978. *The Ticket Book.* The Ticket Book, INC.

Eterno, John A, and Eli B Silverman. 2012. *The Crime Numbers Game: Management by Manipulation.* CRC Press.

Fahim, Kareem. 2006. "Police in Brooklyn Used Illegal Ticket Quotas, Arbitrator Decides." http://www.nytimes.com/2006/01/20/nyregion/police-in-brooklyn-used-illegal-ticket-quotas-arbitrator-decides.html.

Gelman, Andrew. 2005. "Analysis of Variance—why It Is More Important Than Ever." *The Annals of Statistics* 33 (1). Institute of Mathematical Statistics: 1–53.

Gelman, Andrew, and Jennifer Hill. 2006. *Data Analysis Using Regression and Multilevel/Hierarchical Models.* Cambridge University Press.

Kharratzadeh, Milad. 2017. "Splines in Stan." *GitHub Repository.* https://github.com/milkha/Splines_in_Stan; GitHub.

Rayman, Graham A. 2013. *The Nypd Tapes: A Shocking Story of Cops, Cover-Ups, and Courage.* St.

Martin's Press.

Scheidlin, Shira. 2013. "Opinion and Order 08 Civ. 1034 (Sas)." https://ccrjustice.org/sites/default/files/assets/files/Floyd-Remedy-Opinion-8-12-13.pdf.

Stan Development Team. 2016. *RStan: The R Interface to Stan* (version 2.14.1). http://mc-stan.org.

Wallace, Mike. 2017. *Greater Gotham: A History of New York City from 1898 to 1919.* Oxford University Press.