



HCEO WORKING PAPER SERIES

Working Paper



HUMAN CAPITAL AND
ECONOMIC OPPORTUNITY
GLOBAL WORKING GROUP

The University of Chicago
1126 E. 59th Street Box 107
Chicago IL 60637

www.hceconomics.org

Now You See Me, Now You Don't: The Geography of Police Stops

Jessie J. Kalinowski, Matthew B. Ross, Stephen L. Ross*

Abstract

This paper uses state police stop data in Texas to assess patrol activity. We find that both the types of stops and the allocation of resources over space change in darkness relative to daylight, and that the changes in stop type and manpower allocation are correlated within police officers. We also find that the counties receiving more police resources in darkness have a higher share of minority residents. Veil of Darkness (VOD) tests of racial discrimination in traffic stops require that the distribution of motorists be independent of darkness, which is unlikely to be the case without detailed geographic controls.

JEL Codes: K14, K42, J15, H11

Key Words: Police, Traffic Stops, Patrol Locations, Veil of Darkness, Racial Profiling, Racial Discrimination

* Kalinowski: Department of Economics, Quinnipiac University, Hamden, CT (email: jesse.kalinowski@gmail.com); Ross: Wagner School of Public Service, New York University, New York City, NY (email: mbr1@nyu.edu); Ross: Department of Economics, University of Connecticut, Storrs, CT (email: stephen.l.ross@uconn.edu). Acknowledgments: We thank Sharad Goel and the Stanford Open Policing Project for providing us access to the raw Texas Highway Patrol stop data. We are also grateful to Ken Barone and James Fazzalero of the Connecticut Racial Profiling Prohibition Project for their invaluable perspective on policing. All remaining errors are our own.

Now You See Me, Now You Don't: The Geography of Police Stops

The high share of minorities involved in both traffic stops and searches is frequently cited by advocates as representing clear evidence of racial profiling on the part of law enforcement. On the other hand, police often attribute these disparities to differences in guilt over race and ethnicity. The problem in assessing the racial share of traffic stops is that we do not know the composition of motorists at risk of being stopped by police. A recent solution to this problem, the Veil of Darkness (VOD) test, was proposed by Grogger and Ridgeway (2006) and more recently applied by Ridgeway (2009) and Horace and Rohlin (2016). In the VOD test, the racial composition of stops after sunset is used to assess the distribution of stops in daylight (at the same time of day and day of week) since race is less likely to be observed by police after sunset.

A key maintained hypothesis of this VOD approach is that neither motorists nor police change their behavior based solely on the amount of light. However, Kalinowski, Ross and Ross (2016) show that black motorists appear to drive slower during the day when their race can be observed. Further, Fazzalero et al. (2017; 2018) present a series of maps and describe conversations with police departments that suggest that suburban municipal police in Connecticut change patrol locations at night, moving their attention closer to urban centers and shifting away from speeding and towards equipment violation enforcement.

This paper uses state police stop data in Texas to assess patrol activity. We find evidence that both the types of stops and allocation of resources change in darkness relative to daylight, and that the changes in stop type and manpower allocation tend to be correlated within police officers. Further, we find evidence that the counties receiving more police resources in darkness have a higher share of minority residents.

I. Texas Police Stop Data

The paper uses data collected as part of the Stanford Open Policing Project which contains 13.5 million stops made by 3,606 Texas Highway Patrol officers from 2010 to 2015.¹ These officers are assigned to one of nineteen highway patrol districts, and each district contains between 3 and 30 counties, on average approximately 13 counties per district.² The data identifies the location of the stop, date and time of the stop, the reason for the stop i.e. either the type of warning or citation issued, the race and ethnicity of the motorist stopped, and an identifier for the police officer making the stop.³

We establish an inter-twilight window using data from the United States Naval Observatory (USNO) such that the lower bound is the earliest time of day that sunset begins during the year in the easternmost county of the state and the upper bound is the latest time of the end to the evening, civil twilight in the westernmost county. We select as a sample all stops that fall within the inter-twilight window except for those stops that fall during daily civil twilight for the date of the stop, again using the earliest start and latest end of twilight in Texas.

We restrict our sample to officers whose duties appear to be related to patrol and issuing citations by selecting only those officers in the 75th percentile or higher in terms of number of warnings and citations issued within their patrol district.⁴ Then, in order to be conservative, we only use warnings and citations issued by those officers to non-Hispanic white motorists.

¹ As discussed subsequently, our analysis makes use of a restrictive sample of stops made of only white motorists by officers within the 75th percentile of total district stops and having occurred in the inter-twilight sample. We exclude stops made in counties that border the time zone border, i.e. Hudspeth and El Paso counties. There were 1,311,191 total stops in our main analytical sample and 997 total officers.

² See appendix for a map of Highway Patrol Districts.

³ The raw data contains a patrol district indicator for each officer which is one of the few variables in the data with poor coverage. Rather than rely on this indicator, we assign officers to patrol districts based on the locations where they made the majority of their stops within a given month.

⁴ All results robust to selecting officers above the median or above the 90th percentile. See appendix for a Lorenz curve illustrating the distribution of stops over officers.

II. Type of Police Stops

We first document that the distribution of stops by violation type varies considerably between daylight and darkness. Figure 1 shows relatively dramatic shifts in the composition of stops with darkness giving rise to a substantially larger number of stops associated with warnings for equipment violations and daylight being associated with a much larger number of speeding stops, as well as other miscellaneous warnings. These findings are consistent with the anecdotal stories provided by police in Connecticut in Fazzalaro et al. (2017, 2018).

In order to assess whether this variation is systematic, we calculate the fraction of stops separately for daylight and darkness for each of 20 violation types treating warnings or citations as distinct types. We then sum the absolute value of the daylight-darkness difference over all types and divide by the number of types. The mean daylight-darkness difference is 1.9 percentage points. For comparison purposes, random allocation across the 20 types implies an average citation share of 5 percent.

We then resample police officers retaining the county and the time of each officer's stop and selecting a date (the seasonal variation) randomly allowing us to simulate the daylight/darkness treatment for each stop under the null hypothesis of no change between daylight and darkness. The average statistic under the null of no systematic changes is 1.12 with a standard deviation of 0.01, and the fraction of simulations that exceed 1.9 percentage points is far below 1 in a 1,000. Similar results arise restricting our sample to stops made within 28 days on either side of the Daylight Savings Time (DST) following Ridgeway (2009).⁵

⁵ See appendix.

III. Location of Stops

Next, we attempt to describe changes in the allocation of police manpower across counties between daylight and darkness. First, we restrict our sample to officers whose duties appear to be related to patrol and issuing citations by selecting only those officers in the 75th percentile or higher in terms of number of warnings and citations issued within their patrol district, again using only warnings and citations issued to non-Hispanic white motorists. Then, for each officer, we allocate their time hour by hour and day by day based on the county in which they issued one or more warning or citation. If officers issue a warning or citation in more than one county in a given hour, then their time is evenly divided between the counties. The logic behind this approach is to capture the location of officers at a given point in time, without using information on the number of stops, which might be correlated with either darkness or the representation of minority motorists.

We use this allocation to calculate the fraction of total police resources allocated to each county in the state separately by daylight and darkness and then calculate the difference between these two fractions. The 75th percentile of counties has a 0.044 percentage point increase in the share of resources from daylight to darkness, while the 25th percentile has a 0.036 percentage point decrease. For comparison, equal allocation across all counties in the state would imply 0.39 percent of resources allocated to each county so that the 25th and 75th percentile counties both represent a 10 percent shift relative to the average allocation of resources.⁶

Figure 2 shows the results of an example of this exercise for two state police districts located in and around Dallas-Ft. Worth. The figure shades counties based on the differences between share of police resources allocated to each county in daylight and darkness. The orange

⁶ Again, note that we are using 252 rather than 254 counties because we drop the two counties on the Mountain time zone border.

shading shows a substantial shift in resources towards Dallas in daylight and towards the suburban counties immediately surrounding Dallas and Ft. Worth at night. Note that the orange shaded county at the top of the picture in district 1C is the county that contains Sherman, TX, the primary urban area within that district. Unlike the case of Connecticut towns described by Fazzalaro et al. (2017, 2018), the allocation of state police resources in the Dallas-Ft. Worth area appears to be shifted towards suburban and rural counties after sunset.⁷

More generally, we can summarize the extent of change across counties within district by averaging the absolute value of this difference over all counties in Texas. The resulting statistic summed across counties is 0.07 percentage points, which is approximately 18 percent of the mean allocation. Then, we resample police officers from the data retaining the county and the time of each officer's stop and selecting a date (the seasonal variation) randomly allowing us to simulate the daylight/darkness treatment for each stop under the null hypothesis of no geographic differences by daylight. The average statistic under the null of no geographic shift is 0.059 with a standard deviation of 0.001, and the fraction of simulations that exceed 0.07 percentage points is far below 1 in 1,000. All results above are robust to examining only speeding stops, or restricting ourselves to stops made within 28 days of Daylight Savings Time (DST).⁸

IV. Geography of Stops, Citation Type and Racial Composition

Finally, we compare the extent of geographic shifts across counties to the changing pattern over types of stops, the racial composition of counties and the size of the counties. In order to do this, we create a sample of police officer by county observations. For each officer and county, we calculate the daylight to darkness shift in the officer's resource allocation to that county using

⁷ The appendix presents the statewide distribution of daylight to darkness shifts for the annual and the daylight savings time samples.

⁸ See appendix.

either all stops of non-Hispanic whites made by the officer or speeding stops only. For each officer, we also calculate the change in the violation type as the average over all types of the absolute value of the daylight to darkness change in percent of stops of a given type. For each county, we use 2017 census data to calculate the fraction of residents in each county who are black, the fraction of residents who are non-Hispanic white and identify the highest population county in each district.

These results are shown in Table 1. Column 1 contains estimates from regressing officer average absolute change in violation type within district on the absolute value of the resource allocation change by county by officer. Columns 2, 3 and 4 contain estimates from regressing the county racial composition or size indicator on the signed value of the resource allocation change. All four models include district fixed effects and cluster standards errors at the district level. Panel 1 presents estimates for the annual sample and panel 2 for the DST sample.

With one exception, we find a strong relationship between violation type and resource allocation. Officers who substantially reallocate time across counties between daylight and darkness also substantially change the type of stops they make. The standard deviation of the geographic change between day and night is 3.5 percentage points, relative to an average of about 13.5 percent given that the average police officer makes stops in 7.4 counties. Focusing on the three significant estimates, one standard deviation geographic change in policing is associated with between a 0.12 and 0.27 standard deviation changes in the composition of stops over violation type.

In the annual sample, we also find a relationship between county composition and resource allocation changes with counties that contain more minority residents receiving more police resources in darkness. However, the effect is small with a one standard deviation change in the geography of stops associated with composition changes of between 0.2 and 0.3 percentage points.

For speeding stops only, we find evidence that more resources are allocated to the largest population counties. A one standard deviation change in geographic re-allocation of speeding stops is associated with a two percentage point change in the likelihood of a stop's county being the largest population county in the district, relative to a 1 in 13 chance on average.

Neither the demographic nor county size results are robust to the DST sample, but our ability to identify robust demographic patterns may be limited by our reliance on county, as opposed to neighborhoods or cities and towns, as the geographic area. The data does contain additional information on the type of road, and we use this data to run an analysis that is similar to the analysis in Table 1 using the change in officer allocation by county by road type, except that we must use the racial composition of the stops.⁹ Table 2 presents these results. The correlation between the geographic shift in policing and violation type is robust in general, and stronger in magnitude for the geographic changes based on all types of stops. The positive correlation between the geographic shift from daylight to darkness and share of stops that are black is also robust, and now is also significant for the DST sample.¹⁰ The non-Hispanic white estimates are unstable for geography based on all stops across the annual and DST samples, but robust and comparable in magnitude to the race estimates for speeding stops only.

V. Implications for Veil of Darkness Tests for Discrimination

In this paper, we demonstrate that when darkness falls, regardless of the time of day, police officers change their patrolling behavior. The types of stops change substantially at night with a decrease in the number of speeding stops and an increase in number of equipment warnings in

⁹ The county level analysis using share of stops rather than share of population in county generates results that are similar to the findings in Table 1.

¹⁰ The standard deviation of the geography variable by county by road type is 2.9 versus 3.5 for the county variable. As a result, the standardized effect size for the annual sample is larger than the effects in Table 1, and the effects for the DST sample are much larger than the DST estimates from Table 1.

darkness. We also observe a substantial geographic shift between daylight and darkness in where stops are being made, and this geographic shift appears to be linked to the changes in type of stop. We also find some evidence of more stops in darkness in minority counties and in the largest counties. These findings mirror evidence in Connecticut.

Grogger and Ridgeway (2006) focus on a sample of only speeding stops based on concerns that discrimination might be masked because police are focusing at night on different types of stops. The evidence in this paper supports their decision to separate different types of stops when conducting Veil of Darkness (VOD) styles tests. However, the nature of the speeding stops could still change between daylight and darkness. At night, if police are not focused on speeding stops or other moving violations, they may only pull over the most extreme moving violations, and the resulting sample may not be comparable to the sample of stops in daylight when police are focused on routine speeding stops.

Grogger and Ridgeway (2006) also note that seasonal changes in tourist areas might confound identification in the annual VOD sample, but suggest that estimation using subsamples within a window surrounding the DST change or in urban areas would likely mitigate this problem. Our findings suggest that changes in the geographic distribution of stops occur in urban areas and persist through more rigorous sample restrictions.

In practice, researchers need to understand how the geographic distribution of stops changes between daylight and darkness, and then include controls for geography at the appropriate geographic level. Otherwise, VOD tests may fail to find discrimination if police shift enforcement activities in darkness towards more urban areas and/or towards places with greater minority residential share. Simple controls for state police barracks or town/city police departments are

insufficient to address this concern because police appear to be reallocating activity within their own jurisdiction when darkness falls.

References

Fazzalaro, James, Kenneth Barone, Matthew B. Ross, and Jesse Kalinowski. 2017. State of Connecticut Traffic Stop Data Analysis and Findings, 2014-15 supplemental report. Connecticut Racial Profiling Prohibition Project.

Fazzalaro, James, Kenneth Barone, Matthew B. Ross, and Jesse Kalinowski. 2018. State of Connecticut Traffic Stop Data Analysis and Findings, 2015-16 supplemental report. Connecticut Racial Profiling Prohibition Project.

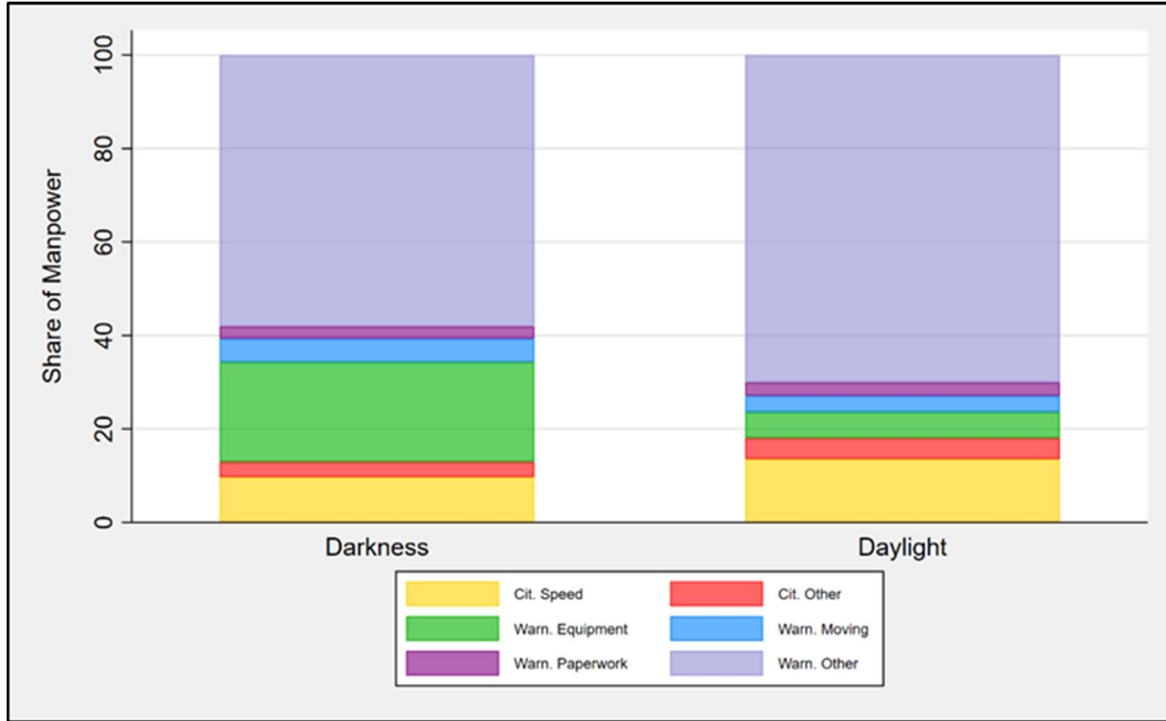
Grogger, Jeffrey and Greg Ridgeway. 2006. "Testing for Racial Profiling in Traffic Stops from Behind a Veil of Darkness". *Journal of American Statistical Association*, 101, 878-887.

Horrace, William C. and Shawn M. Rohlin. 2016. "How Dark Is Dark? Bright Lights, Big City, Racial Profiling." *Review of Economics and Statistics* 98, 226-232.

Kalinowski, Jesse, Matthew Ross and Stephen L. Ross. 2017. "Endogenous Driving Behavior in Veil of Darkness Tests for Racial Profiling". *Human Capital and Economic Opportunity (HCEO) Working Paper* 2017-17.

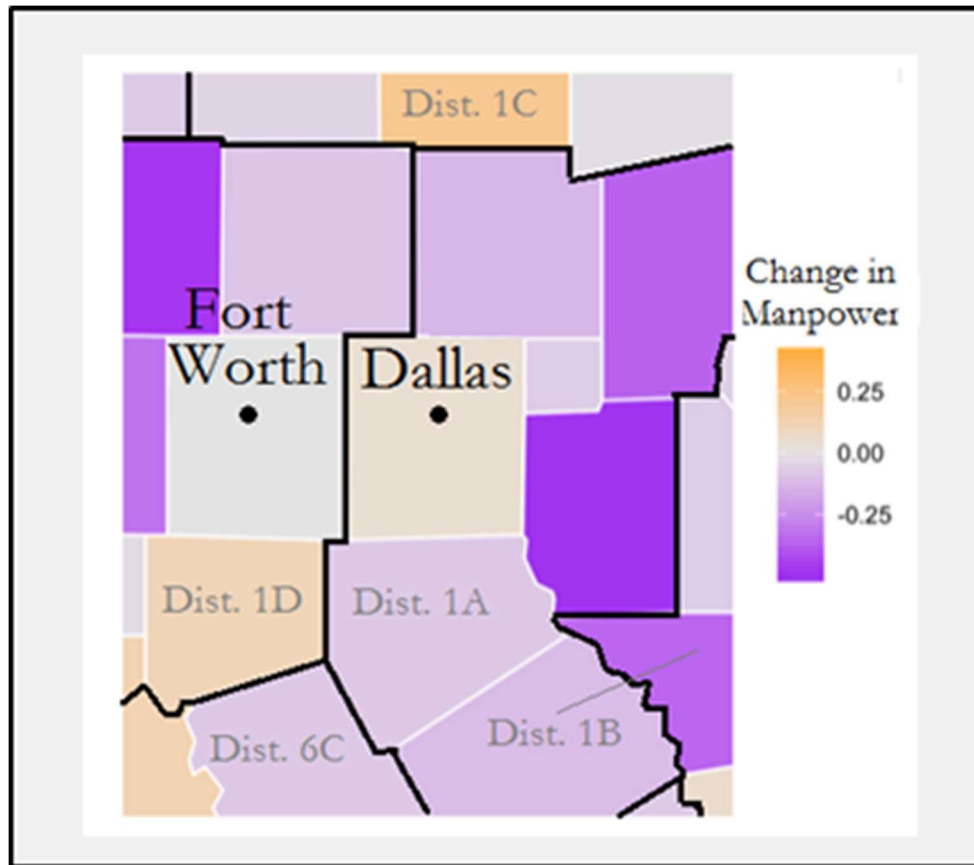
Ridgeway, Greg. 2009. "Cincinnati Police Department Traffic Stops: Applying RAND's framework to Analyze Racial Disparities". *Rand Corporation: Safety and Justice Program*.

Figure 1. Change in Violation Distribution



Notes: Each column shows the fraction of police stops of non-Hispanic white motorists associated with each type of citation or warning. Stops with multiple citations or warnings have their contribution divided evenly between each category.

Figure 2. Change in Manpower



Notes: Daylight to darkness change in share of state police resources based on non-Hispanic white stops. Orange shading represents a larger fraction of resources allocated in daylight, with purple representing more resources in darkness.

Table 1. Policing Changes by County

Outcome	Violation Type	Black	Non-Hispanic White	Largest County
Annual Sample				
All Stops	0.035*** (0.008)	0.050*** (0.010)	-0.055* (0.031)	-0.001 (0.001)
Speeding Stops Only	0.000 (0.003)	0.043*** (0.014)	-0.082*** (0.015)	0.006** (0.002)
Daylight Savings Time Sample				
All Stops	0.076*** (0.01)	0.007 (0.02)	0.018 (0.045)	-0.001 (0.001)
Speeding Stops Only	0.046*** (0.011)	0.003 (0.012)	0.008 (0.023)	-0.000 (0.001)

Notes: Independent variable is percentage changes in police resources in county between daylight and darkness with percentages scaled from 0 to 100: absolute value for violation changes and actual signed changes for county demographics and the largest county dummy variable. Dependent variables are listed in the top row. Violation type changes are standardized, and percent black and non-Hispanic white range from 0 to 100. Models include district fixed effects and standard errors are clustered by district. *** 1 percent, ** 5 percent and * 10 percent significance level.

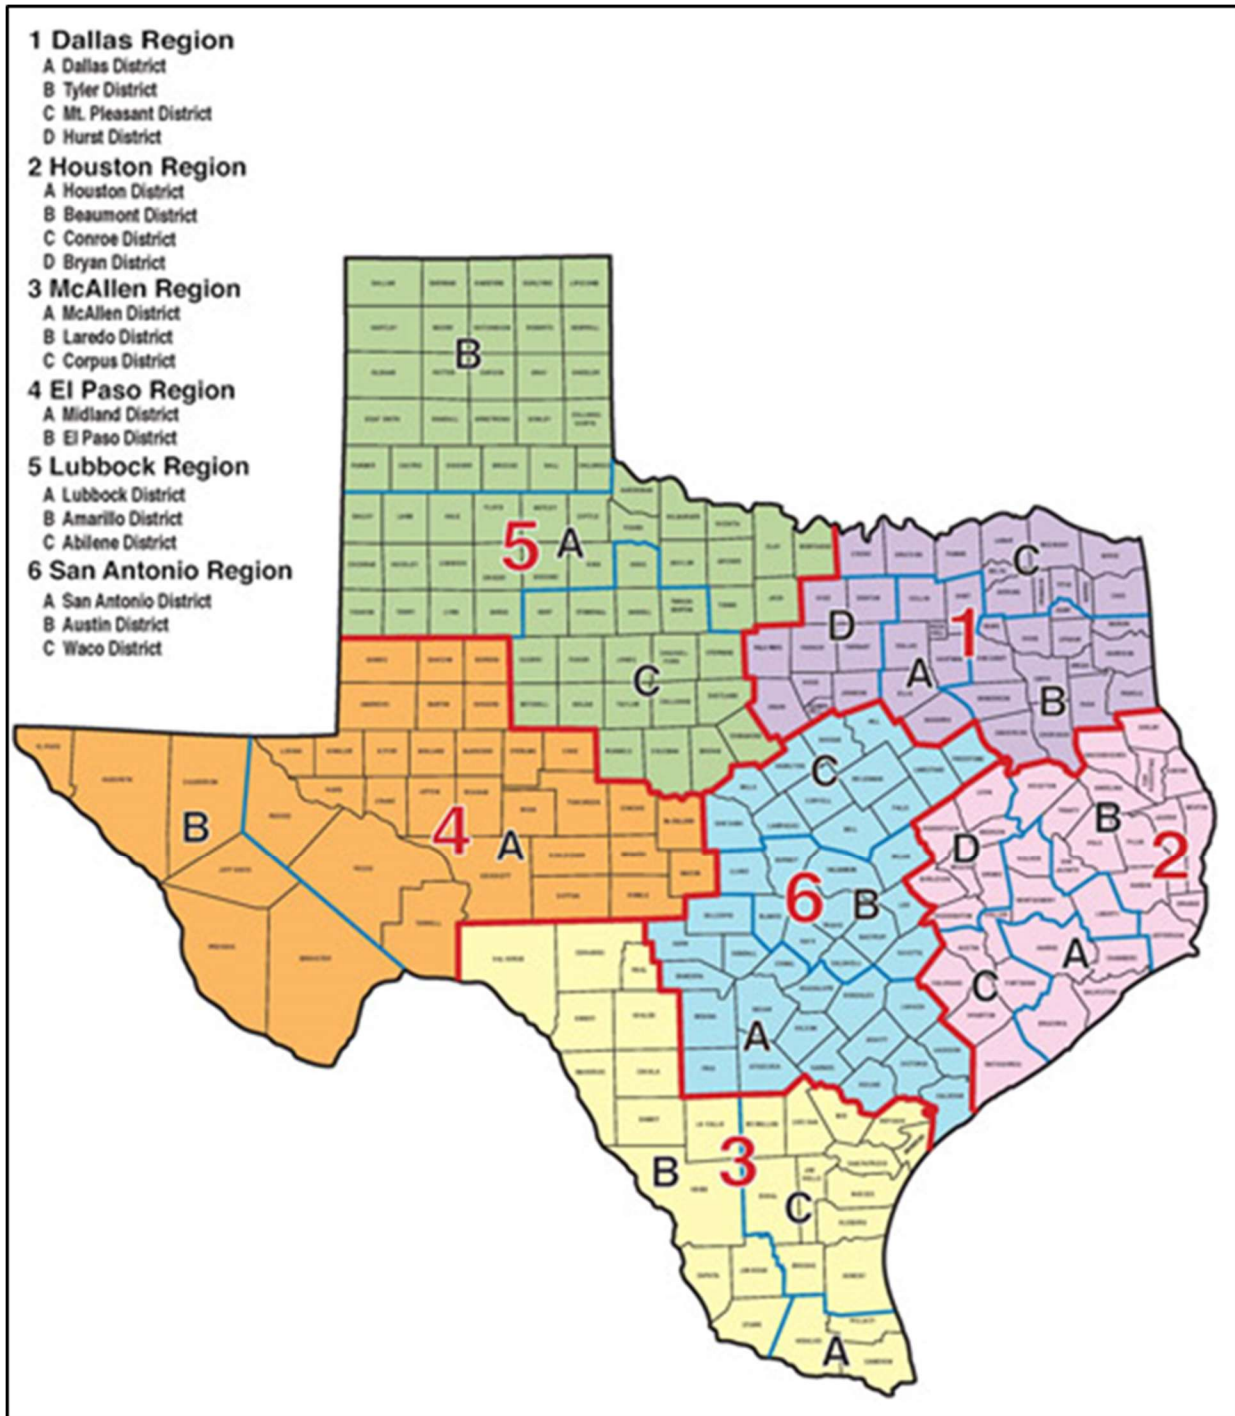
Table 2. Policing Changes by County by Road Type

Outcome	Violation Type	Black	Non-Hispanic White
Annual Sample			
All Stops	0.066** (0.024)	0.124* (0.061)	0.386*** (0.081)
Speeding Stops Only	-0.000 (0.004)	0.060*** (0.014)	-0.083*** (0.021)
Daylight Savings Time Sample			
All Stops	0.146*** (0.032)	0.065* (0.036)	-0.006 (0.079)
Speeding Stops Only	0.037*** (0.009)	0.023*** (0.007)	-0.036** (0.014)

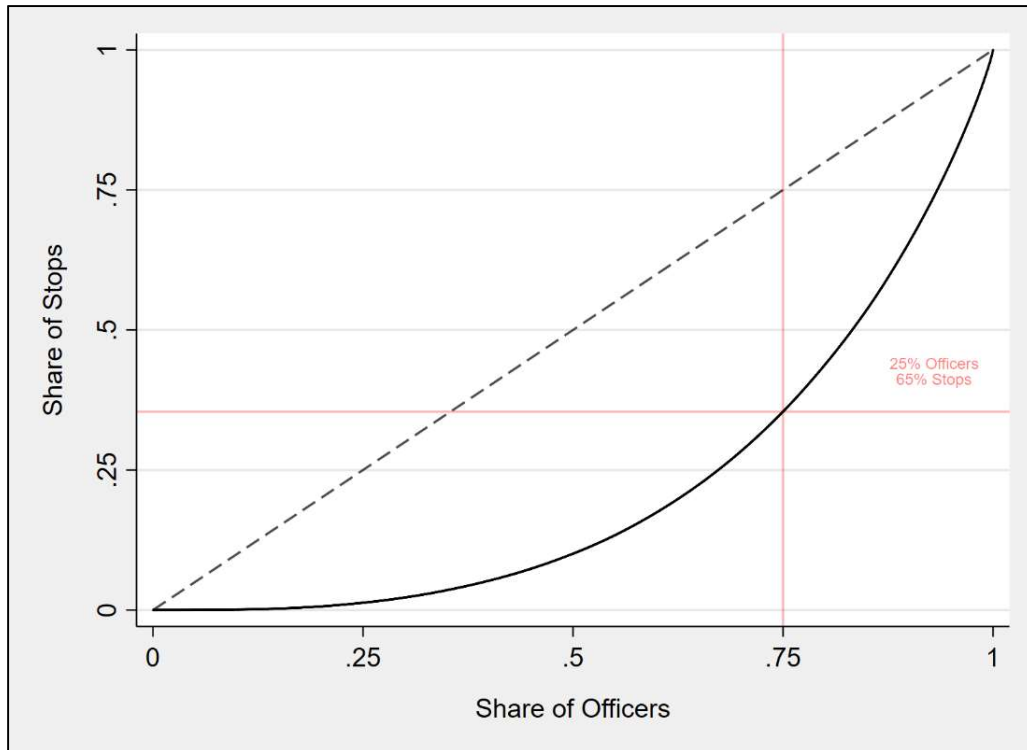
Notes: Independent variable is percentage changes in police resources in county by road type cell between daylight and darkness with percentages scaled from 0 to 100: absolute value for violation changes and actual signed changes for racial share of stops in county by road type cell. Dependent variables are listed in the top row. Violation type changes are standardized, and percent black and non-Hispanic white stops range from 0 to 100. Models include district fixed effects and standard errors are clustered by district. *** 1 percent, ** 5 percent and * 10 percent significance level.

Appendix

Appendix Figure 1. Map of Texas Highway Patrol Districts and Counties

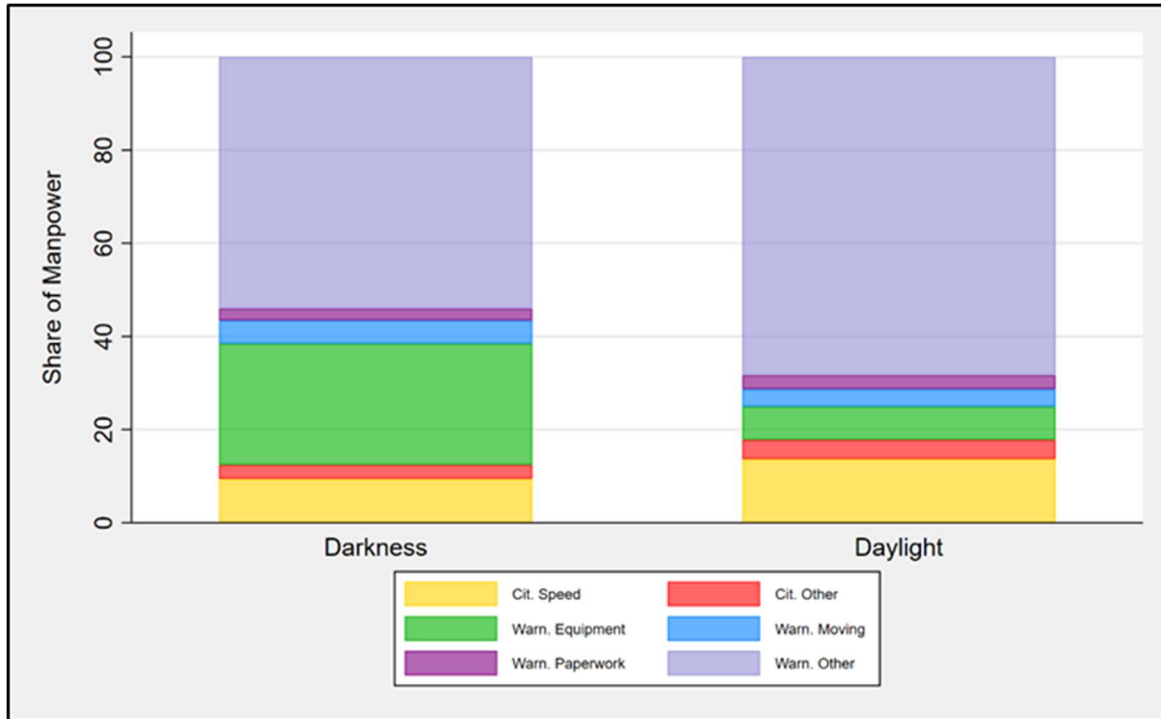


Appendix Figure 2. Distribution of Officers and Stops, Annual Inter-Twilight Sample



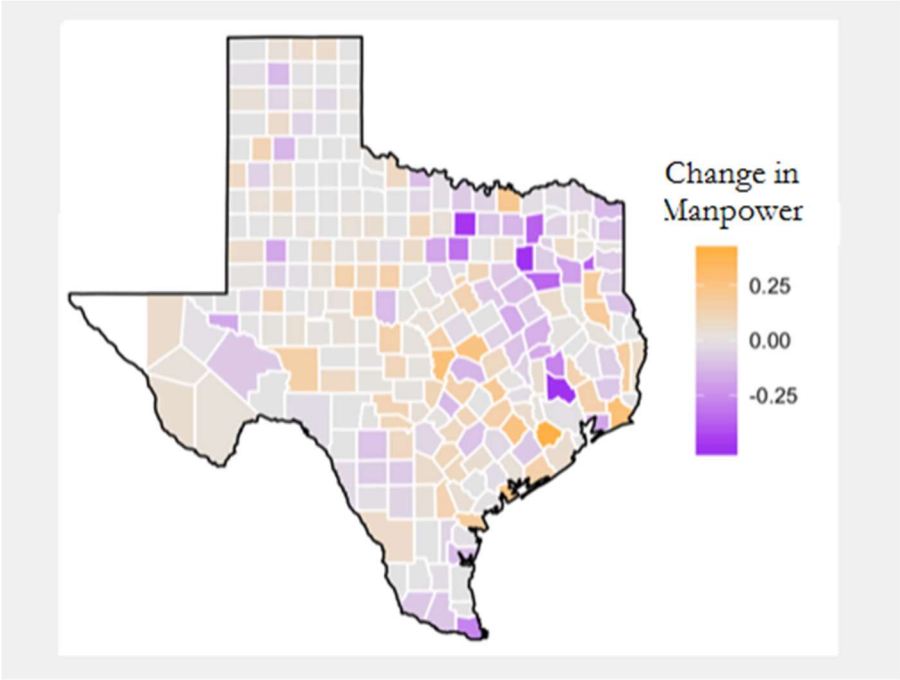
Note: Officers were selected from the 75th percentile by individual patrol district. For descriptive purposes, this figure presents the distribution across all patrol districts. All findings are robust to selecting officers at various other thresholds including the 50th and 90th percentile. The sample includes all officers who made any stops in the annual inter-twilight window.

Appendix Figure 3 Distribution of Types of Stops for Daylight Savings Time Sample



Appendix Figure 4. Statewide Map of Geographic Shift in Manpower

Panel 1. Annual Sample



Panel 2. Daylight Savings Time Sample

