City University of New York (CUNY)

CUNY Academic Works

Economics Working Papers

CUNY Academic Works

2016

Speeding, Punishment, and Recidivism – Evidence from a Regression Discontinuity Design

Markus Gehrsitz

CUNY Graduate Center, CUNY City College

Follow this and additional works at: https://academicworks.cuny.edu/gc_econ_wp

Part of the [Economics Commons](#)

[How does access to this work benefit you? Let us know!](#)

More information about this work at: https://academicworks.cuny.edu/gc_econ_wp/11

Discover additional works at: <https://academicworks.cuny.edu>

This work is made publicly available by the City University of New York (CUNY).

Contact: AcademicWorks@cuny.edu

**Speeding, Punishment, and Recidivism –
Evidence from a Regression Discontinuity Design**

Markus Gehrsitz

Working Paper 11

Ph.D. Program in Economics
CUNY Graduate Center
365 Fifth Avenue
New York, NY 10016
March 2016

I thank the German Federal Motor Transport Authority (KBA) for granting me on-site access to the main data source for this article. I am grateful to Dr. Dirk Hillebrandt for his patience and his dedication to this project. I would like to thank Michael Winkler for regularly enlightening me with respect to the idiosyncrasies of the German traffic penalty catalogue. I benefited from the helpful comments and suggestions of Michael Grossman and David Jaeger. This work received support through the CUNY Doctoral Student Research Grant (DSRG) Program.

© 2016 by Markus Gehrsitz. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Speeding, Punishment, and Recidivism –
Evidence from a Regression Discontinuity Design
Markus Gehrsitz
JEL No: K42, I12, I18

ABSTRACT

This paper estimates the effects of temporary driver's license suspensions on driving behavior. A little known rule in the German traffic penalty catalogue maintains that drivers who commit a series of speeding transgressions within 365 days should have their license suspended for one month. My fuzzy regression discontinuity design exploits the quasi-random assignment of license suspensions caused by the 365-day cutoff and shows that 1-month license suspensions lower the probability of recidivating within a year by 20 percent. This effect is not driven by incapacitation and indicates that temporary license suspensions are an effective tool in preventing traffic transgressions.

Markus Gehrsitz
NBER
5 Hanover Square, Suite 1602
New York, NY 10004
mgehrsitz@gradcenter.cuny.edu

He who undertakes to punish with reason does not avenge himself for the past offense, since he cannot make what was done as though it had not come to pass; he looks rather to the future, and aims at preventing that particular person and others who see him punished from doing wrong again. (Plato, Protagoras p. 324)

1 Introduction

Worldwide, more than 1.2 million people die every year in traffic accidents (WHO, 2013). In 2010, traffic injuries in the US claimed the lives of 45,342 Americans and caused medical and work loss costs of more than \$100 billion. Motor vehicle crashes are the leading cause of death for those aged 15 to 35 (CDC, 2015). Numerous laws and regulations, such as mandatory seatbelt use, speed limits, and blood alcohol concentration (BAC) limits, have been enacted in order to promote road safety. Most developed countries have an elaborate penalty catalogue in place that aims to punish and deter traffic offenders. Temporary license suspensions are a crucial component of these penalty catalogues. This measure not only incapacitates traffic offenders by taking them off the streets for a short period of time. It is also supposed to provide a “shot across the bows” by prompting offenders, who temporarily experience the inconveniences of life without a driver’s license, to change their ways in the long-run and drive more responsibly once they get their license back.

Yet, little is known about the effectiveness of temporary license suspensions. Economic theory provides little guidance on this issue. On the one hand, Gary Becker’s (1974) model of the “rational criminal” predicts that temporary license suspensions should have no long-run effects. In this classical model, criminals weigh the expected costs of committing a crime against the expected benefits. As long as a temporary license suspension has no effect on the degree of a penalty for future transgression (e.g. through a three strike law), the measure should not change this cost-benefit-trade-off. Traffic offenders are therefore just as likely to commit another transgression before a temporary license suspension as they are after they get their license back. On the other hand, the behavioral approach to the economics of crime (Jolls et al., 1998, among others) offers many channels through which a temporary

license suspension might indeed have an effect. For instance, offenders might update their beliefs about the actual cost of punishments or the probability that an offense is detected.

This is therefore ultimately an empirical issue. As Levitt and Miles (2007) point out, the main challenge for empiricists is to distinguish causation from correlation. A naive comparison of the recidivism behavior of offenders who were punished for their criminal acts with those who were not, cannot shed much light on the question of how effective penalties are in deterring future crimes. After all, these penalties are not randomly assigned. People who get punished because they committed a crime might be intrinsically prone to committing crimes. Comparing their rates of recidivism to those of law abiding citizens is thus an apples to oranges comparison. Traffic transgressions are no exception to this problem. As shown below, a “naive” OLS regression yields a *positive* effect of license suspension on recidivism even when conditioning on age, sex, and state fixed effects.

This illustrates the challenges in answering a seemingly simple question. The goal of this paper is to overcome these challenges. For that purpose, I exploit a little known but widely applied rule in the German penalty catalogue for traffic violations which leads to a quasi-random suspension of some individuals’ driver’s licenses. Specifically, this rule maintains that a person who commits two major speeding violations within 365 days, should have her license revoked for one month. This gives rise to a fuzzy regression discontinuity design where those to the left of the 365-day cutoff are likely to have their license revoked while those to the right of the cutoff retain theirs. The rule appears to be obscure and complex enough to prevent sorting to the right of the threshold. At the same time, the rule is very much enforced by the traffic authorities. I find that a loss of license for one month reduces the probability of recidivating within a year by about 20 percent. This result is robust to the choice of different time windows and estimation techniques.

2 Background and Institutional Framework

Endogeneity issues in the relationship between punishment and criminal activity have been recognized since at least the late 1990s (Levitt, 1996). Ever since, economists have come up with various strategies to exploit some source of exogenous variation in order to isolate causal effects of punishment on recidivism. Kessler and Levitt (1999) exploit sentence enhancements that are exogenously induced by California’s Proposition 8 to evaluate the effect of harsher sentences on crime. They find that harsher punishments reduce crime substantially. Helland and Tabarrok (2007) utilize idiosyncrasies in the same state’s three strike policy to isolate a large and significant deterrence effect of the policy. Lee and McCrary (2009) use the fact that juvenile criminals tend to be sentenced as adults once they turn eighteen, and find only small deterrence effects of more severe penalties.

Drago et al. (2009) analyze sentence reductions for certain parts of the Italian prison population due to a collective clemency bill. Their results suggest that a reduction in prison sentences actually reduced recidivism. Di Tella and Schargrodsky (2013) use Argentinian judges’ ideological differences as an instrument in evaluating the effectiveness of electronic monitoring compared to incarceration. They find that this more lenient treatment reduces recidivism. Green and Winik (2010) and Aizer and Doyle (2015) exploit the random assignment of criminals to judges with different incarceration tendencies (an identification strategy pioneered by Kling (2006)) and find that juvenile imprisonment increases recidivism probabilities and has adverse effects on human capital accumulation. In the arena of traffic violations, DeAngelo and Hansen (2014) show that a mass layoff of Oregon State Police in 2003 was associated with large increases in accidents and traffic fatalities. In a separate study, Hansen (2015) exploits blood alcohol content (BAC) cutoffs that determine whether a transgression is treated as a driving under the influence (DUI) offense, or even an aggravated DUI offense. His regression discontinuity results suggest that getting caught behind the wheel with a BAC just above the DUI threshold reduces recidivism by 17 percent.

My study is similar in spirit in that it exploits an exogenous increase in the severity of the penalty for a speeding offense to isolate the causal effect of a 1-month license sus-

pension on repeat offenders. The source of this exogenous variation is an idiosyncrasy in the German traffic penalty catalogue. This catalogue generally provides for three different types of penalties: money fines, (demerit) points entered into one's central traffic registry account, and license suspensions. The degree of the penalty is determined by the seriousness of the offense. Table 1 provides an excerpt from the penalty catalogue, specifically for speeding offenses. Small transgressions, such as driving 10 km/h over the stipulated speed limit, are fined with small financial penalties. For more severe transgressions, points in the central traffic registry are added. The point system in Germany is similar to those in place in most EU countries and several US states. Points received from different transgressions and different types of transgressions (speeding, DUIs, running red lights, etc.) accumulate. If a person does not commit a transgression for two years, then all points are erased. If a person does commit a transgression within two years, however, new points are added to the existing stock and the two year expungement period starts afresh. An offender permanently loses her license once her stock of points rises to 18.¹

Finally, for severe transgressions, temporary license suspensions are handed out on top of points and money fines. For instance, a person speeding 45 km/h over the limit will have her license suspended for 1 month. The penalty catalogue also distinguishes between offenses that took place on highways and offenses that took place in built-up areas, e.g. residential neighborhoods. Fines differ in severity for different types of transgressions ranging from aggravated DUIs to driving without appropriate snow chains.²

The road traffic law (BKatV), which constitutes the legal basis for the penalty catalogue, also has multiple additional provisions. The legal text is currently in its 87th edition, has 3290 pages, and weighs about five pounds. One additional provision is for "persistent delinquencies." It maintains that "a temporary license suspension [of usually 1 month] shall ordinarily be handed out if the operator of a motor vehicle commits a speeding transgression

¹The point system was reformed in May 2014; all information therefore refers to the law prior to this reform.

²The current version of the German penalty catalogue for traffic offenses can be found at <http://www.kba.de>

of at least 26 km/h within 1 year after another speeding transgression of at least 26km/h has been committed and the corresponding penalty has obtained legal force” (§4 Abs. 2 Satz 2 BKatV; own translation). This provision will henceforth be referred to as the “365-day rule.” For instance, a person who within a few months is caught twice exceeding the speed limit on a highway by 28 km/h falls under the 365-day rule and will have her license temporarily revoked even though, according to the penalty catalogue, she should on aggregate only be fined €160 and receive 6 demerit points (see Table 1). Note that the wording of the provision is tricky. If you read the above text carefully, you will notice that the period, in which an offender is at risk of falling under the 365-day rule, only begins after the penalty for the first transgression “has obtained legal force.” The difference between the date of the transgression and the date on which the associated penalty obtains legal force might seem minor at first yet the median difference between those two dates is 66 days. In other words, the day count which determines whether an offender has her license suspended due to the 365-day rule does not start immediately after committing the first transgression but with a substantial time lag. This wrinkle in the law is a big source of confusion among offenders and - as we will see in Section 4 - prevents sorting to the right of the 365-day cutoff. Fortunately, both the date of a transgressions and the date on which the corresponding penalties obtain legal force are recorded in the data, so I can properly ascertain which offenders fall under the rule.

The 365-day rule provides a cutoff that can be exploited in a regression discontinuity setting. For instance, this rule requires that a person who commits her second transgression within 365 days after the penalty for the initial transgression has obtained legal force, should have her licenses suspended for one month. A person who commits her second transgression on day 366 should keep her license. By comparing the recidivism behavior of these two groups of people, who should be very similar except for the degree of the penalty they receive, I can obtain an unbiased estimate of the effect of a 1-month driver’s license suspension.

My identification strategy rests on two assumptions that need to be met to guarantee internally valid estimates. First, drivers by and large must not be aware of this regulation

or, at the very least, they should not gear their driving behavior accordingly. For instance, this assumption would be violated if drivers drove more recklessly once 365 days have passed since the penalty for a previous major speeding transgression obtained legal force. In Section 4, I will test this assumption. The data will support this assumption and will show that sorting to either side of the 365-day cutoff is not common, most likely because the wording of the provision is not straightforward and the rule is somewhat obscure. Second, traffic authorities must enforce this regulation. It is apparent from the above quote that the authorities only “shall ordinarily” hand out licenses suspension and thus have considerable wiggle room. But, Section 4 will also show that the authorities to a great degree adhere to this rule, thus creating a discontinuity in the assignment of drivers to license suspensions.

3 Data

The source of data for this study is the German central traffic registry (“Verkehrszentralregister” or VZR). The VZR is administered and maintained by the Federal Motor Transport Authority (“Kraftfahrt-Bundesamt” or KBA) which is based in Flensburg, Germany. The VZR contains an account with a unique ID for every traffic offender who has committed a transgression that was sanctioned with at least one demerit point. Offenses are usually first recorded by local traffic authorities. These local agencies then transmit information on the date and type of transgression as well as the corresponding penalty to the KBA. In some instances, this information is transmitted by the courts rather than local traffic agencies. The transmissions also contain information about the offender’s sex, age, state of residence, as well as the date on which a penalty obtains legal force.

The VZR is an active registry. Persons who do not commit a traffic transgression for two years not only get their points total set to zero, but are erased entirely from the data base.³ If someone commits a transgression *after* her account has been erased due to this

³This expungement rule was changed during the 2014 reform of the VZR. Now, each offense is associated with a unique expungement period depending on the severity of the offense. In this study only data pre-dating this reform are used.

2-year expungement period, she starts with a clean slate. That is, she receives a new ID and cannot be linked to former VZR entries. The VZR was digitalized in 1999. Offenders who have entries dating back beyond 1999 that have not been expunged yet, are still administered using paper files. For this study only offenders who are administered using digital files (and thus show up in the data base) could be used. While every citizen has a right to enquire his point total, the data in the VZR are usually not made available to outside researchers. What is more, new entries and deletions changes the composition of the VZR everyday. Fortunately, the point system was reformed on 1 May 2014. On this occasion, a dataset containing the digital VZR population as of 30 April 2014 was created. The KBA kindly made an anonymized version of this excerpt available to me for on-site analysis. The data set contains more than 10.5 million entries pertaining to about 6.3 million distinct offenders.

Obviously not all observations in this dataset can be used for this study. For one, a sufficiently large follow-up period is required since my main outcome of interest is recidivism. For example, transgressions that occurred on April 29th 2014 cannot be used for this analysis due to the lack of a follow-up period. More importantly, this study tries to isolate exogenous variation in the severity of punishment to assess the effect of license suspensions on recidivism. This exogenous shock is induced by the above mentioned 365-day rule, so only offenders who have a chance to be affected by this rule should be used for analysis. The following steps describe the selection process that identifies observations that become part of my “experiment.”

First, it should be noted that the 365-day rule only applies to speeding transgressions of 26km/h or more. More precisely, a 1-month license suspension may be imposed on speeders whose offense by itself would not have resulted in a license suspensions. As can be seen from Table 1, there are two cases that qualify for this rule: speeding 26-40 km/h above the limit on highways and speeding 26-30 km/h above the limit in built-up areas. Second, the rule can only affect persons who have previously committed another speeding transgression of 26km/h or more. Throughout this article, I will refer to the earlier of these two transgressions as the “original transgression” and the second transgression will be referred to as the

“treatment transgression.” When I refer to the time difference between the treatment and the original transgression, I am referring to the number of days that have passed between the date on which the penalty for the original transgression obtained legal force and the date of the treatment transgression. In other words, throughout this paper, I account for the wrinkle in the law that delays the start of the day count. Only persons who have both an original transgression and a treatment transgression (i.e. two speeding offenses of 26km/h or more) can enter the final “discontinuity sample”. I also require the time difference between both transgressions to be no more than 545 days and at least 186 days.⁴ In essence, this puts a 180-day window around the 365-day threshold that determines whether a 1-month driver’s license suspension is issued. I will assess the robustness of my results to different time windows around the cutoff in Section 5.

The outcome of interest is recidivism, i.e. the probability of committing yet another (third) offense. This by itself requires a sufficiently large post-treatment time window. This time window has to be even larger since it can take a few months for transgressions to show up in the data. For example, offenders may appeal against prescribed penalties in court. Transmission to the KBA and recording information into the VZR also takes some time. On average, it takes about two months from the date of the actual offense until it shows up in the data base, after 5 months more than 99% of incidents actually show up in the data. My observational period ends on 30 April 2014. Accounting for an at-most 5 month delay and in order to evaluate a 12 months recidivism window, only offenders who have committed their treatment transgression before 1 December 2012 can be used in this analysis.

Another complication arises from the 2-year expungement period. Individuals who do not commit a traffic transgression for two years are erased from the data set. By inversion, individuals who still are in the data base must have committed an offense in the past two years. That is, individuals who committed an offense before 1 May 2012 must necessarily have committed another offense subsequently. For these individuals there will thus be no variation in the outcome. As a result, individuals whose treatment transgression predates

⁴Again, this is accounting for the lag caused by the time that passes between the day of the original transgression and the day on which the penalty for this transgression has obtained legal force.

May 1 2012 also need to be excluded from the discontinuity sample.

Figure 1 illustrates the generation of my final sample. Only observations who have their treatment transgression within the dashed frame (1 May 2012 or later, but before 1 December 2012) can be used for the analysis. Take person A at the top of the graph whose transgressions over time are illustrated by triangles. This person has a transgression within the framed time period, but this is her original transgression, not her treatment transgression. Her (second) treatment transgression is outside the dashed frame, as indicated by a second triangle. Moreover, this person reoffends as there is yet another triangle further to the right. I would most likely fail to observe this recidivism event, however, since it takes place after 1 December 2013 and thus may not yet show up in the data. Person A is therefore excluded from my sample. Person B, on the other hand, will be included in the sample. Her treatment transgression falls into the framed area, giving me a sufficiently large time-window to detect recidivating behavior. In this instance, person B indeed recidivates within 12 months which I observe. Person C is similar and will also be included in my sample. The main difference to person B is that person C does not recidivate. Since I allow for a sufficiently large follow-up period, I can be sure that as a matter of fact there was no further transgression in the year following the treatment transgression. In other words, the lack of another offense for this person is not due to lagged reporting. Finally, person D and everyone else who has their treatment transgression prior to 1 May 2012, by virtue of the expungement period, must necessarily have recidivated. Otherwise, they would no longer be in the sample. That is, there is no variation in the outcome for these observations, leading me to drop person D.

Once these restrictions have been imposed, a “discontinuity sample” emerges. It consists of 31,400 persons. Each person has a treatment transgression that occurred between 1 May 2012 and 30 November 2012; each person also has a original transgression for which the date on which the corresponding penalty has obtained legal force predates the date of the treatment transgression by at least 186 days and at most 545 days. For about half the sample, 365 days or less passed between these two points in time. Borrowing terminology from the potential outcome framework (Angrist et al., 1996), these observations constitute

the “treatment group”. Members of this group fall under the 365-day-rule and should have their licenses suspended for 1 month. The remainder of the discontinuity sample constitute the “control group”. Members of this group have committed similar transgressions but due to the timing of their offenses mostly retain their driver’s license. Persons in both treatment and control group may or may not recidivate, i.e. may or may not have committed a (third) major traffic transgression after their treatment transgressions.

If the descriptive statistics of Table 2 are any indication, it appears as if license suspensions have an effect. Rates of recidivism are 24.4% and 26.5% for the treatment and control group, respectively. These differences are statistically significant at the 1% level (see row 10 and columns (1) and (2) of Table 2). The difference actually tends to increase the more I limit the sample to observations with treatment transgressions closer to the 365-day cutoff. On the other hand, demographic factors that are reported with the data are quite balanced across treatment and control group, indicating a quasi-random separation of the sample. For instance, the average age in the treatment sample is 42.66 years, in the control sample it is 43.02. The null hypothesis that there is no difference in means cannot be rejected at the 1% level. The means for all other covariates are also very similar, and formal t-tests for differences in means fail to reject the null in the vast majority of instances.

I also distinguish between speeding-specific recidivism and general recidivism. The bottom two rows of Table 2 provide the corresponding means and standard deviations. Since 1-month license suspensions are quasi-randomly assigned due to a speeding incidence, one might hypothesize that speeding recidivism is particularly deterred. Therefore a dummy is defined that equals one if another *speeding* transgression occurs after the treatment transgression. The equivalent dummy for general recidivism is equal to one if any transgression (DUI, speeding, running a red light,...) is recorded in the post-treatment period.

4 Methods

The goal of this study is to exploit the exogeneous variation in penalties induced by the 365-day rule. The rule maintains that a 1-month license suspension shall be levied on offenders who - accounting for the wrinkle in the law that delays the start of the day count - commit two major speeding transgressions within 365 days, but not on those who commit two such transgressions within 366 days or more. This will allow for the identification of the causal effect of a temporary license suspension on the probability of reoffending. My identification strategy will only be valid if the 365-day rule is actually applied and results in a discontinuity in the assignment to treatment. Figure 2 illustrates that this is indeed the case. The x-axis shows the running variable, i.e. the number of days that have passed between the date on which the penalty for the original transgression has obtained legal force and the date of the treatment transgression. For each bin, I calculate the fraction of offenders within that bin who have had their license suspended for 1 month. The position of each point relative to the y-axis yields information about these fractions.

If the 365-day rule was strictly applied, everybody to the left of the red vertical line should have her license suspended for 1 month in addition to the prescribed money and point penalties. Everybody to the right or on the line should keep their license and merely suffer the prescribed money and point penalties. Such a “sharp” separation into treatment and control group is not present in this case. There is, however, a big drop in the probability of having one’s license temporary revoked at day 366. To the left of the cutoff around three quarters of offenders lose their license for 1 month and to the right of the cutoff a mere 1.7 percent of offenders are hit with a 1-month license suspension. In other words, there is a huge drop in the probability of having one’s license suspended due to the treatment transgression once 365 days have passed since the penalty for the original transgression obtained legal force. Likely reasons for receiving the treatment on the “wrong” side of the cutoff are involvement in an accident or repeat offending in terms of non-speeding transgressions.⁵ In

⁵Cases in which exactly 366 days have passed between the original transgression and the treatment transgressions have their licenses suspended more frequently than most other cases to the right of the cutoff. A

such instances, the local traffic authorities may mandate a temporary license suspension. The reasons for imperfect compliance on the left hand side are numerous. In instances very close to the cutoff, local traffic authorities or judges may be lenient and choose to not invoke the 365-day rule. In cases of hardship, for example elderly or disabled drivers who have no other means of transportation than their vehicles, local traffic authorities or judges may also be lenient. Similarly, offenders from rural areas, commuters who are dependent on their car, or professional truck drivers might be able to keep their licenses. In general, any penalty notice can be appealed in court and judges may override a suspension if the offender shows remorse or accepts a higher monetary and/or point penalty in lieu of the temporary license suspension.⁶

By and large, traffic authorities follow the 365-day rule which induces a big drop in the probability of having one's license temporarily suspended at the expected threshold. It makes sense, therefore, to take a look at the relationship between the outcomes and the running variable. Figure 3 plots the two recidivism outcomes of interest against the number of days that have passed between the date on which the penalty for the original transgression obtained legal force and the date of the treatment transgression. The running variable is aggregated into 3-day bins. The size of each circle indicates the number of observations in each bin. The position of each circle, relative to the y-axis, indicates the fraction of offenders within a bin who recidivated within 12 months. The recidivism period is extended to 13 months for those who actually have had their license suspended for 1 month in order to account for incapacitation effects as drivers without a license naturally have less of an opportunity to reoffend.

In Figure 3a there is clearly a discontinuity at the 366-day cutoff. This jump yields a first rough estimate of the reduced form (intent-to-treat) effect. The graph suggests that likely reason is that 2012 was a leap year which may have led to confusion among the local traffic authorities as to whether these cases should fall under the 365-day rule.

⁶Note that this kind of selection issue is *not* a threat to my identification strategy. Rather it illustrates the local interpretation of any regression discontinuity coefficient. This issue will be further discussed in Section 7.

offenders who are assigned to treatment (i.e. license suspensions) are three to four percentage points less likely to recidivate than offenders who are not assigned to treatment. Figure 3b focuses on speeding recidivism. Here the discontinuity is less obvious, but there still appears to be a jump of about 2 percentage points at the 366-day threshold. Note that the lowess lines in both figures are merely superimposed to better visualize the pattern in the data but may very well suffer from boundary bias close to the 365-day cutoff. In order to obtain a visual estimate of the size of the treatment effect it is more important to focus on the position of the cloud of points, especially around the cutoff, than to study the position of the lowess lines at the boundary.

Figures 2 and 3 make a compelling case that a) the 365-day rule invokes a quasi-random assignment of license suspensions and b) this assignment indeed has an effect on future recidivism behavior. However, the internal validity of any coefficient obtained through this setup would be in jeopardy if offenders were very much aware of the 365-day cutoff and geared their driving behavior accordingly. Fortunately, there is little indication that this is the case. First, anecdotal evidence suggests that the vast majority of drivers is not even aware that this rule exists. There is a vast amount of online forums in which repeat offenders who fall under the 365-day rule express their shock about their license suspensions. Second, the wrinkle in the provision that starts the day count only after the penalty for the original transgression has obtained legal force, makes it hard for offenders to keep track of whether they are still at risk of falling under the 365-day rule. That is, if they were aware of the exact wording of the law to begin with.

Data back up this claim. If at least some drivers were aware of the 365-day rule and all its wrinkles, and were able to keep track of the exact day count, one would expect “bunching” on the right-hand side of the 365-day cutoff. Drivers would drive more carefully than usual until the 365-day rule no longer applied to them. This would result in a spike of traffic transgressions on days 366-400. Yet Figure 4 gives little indication that this is indeed the case. The frequency of treatment transgressions (in 3-day bins) is very evenly distributed with around 100 transgressions per day on each side of the cutoff. Most notably,

there is no spike in treatment transgressions from offenders whose penalties for their original transgression obtained legal force just a bit more than 365 days ago. One might be mildly concerned about the small drop in the transgression frequency on days 360-362. Yet, the frequency rebounds to above-average levels on days 363-365. Drops in frequency of even greater magnitude can also be observed elsewhere in the distribution. This visual analysis is consistent with McCrary's (2008) more explicit density test for manipulation at the cutoff. The test implies a log difference in height of $-.021$ with a standard error of 0.033 . In other words, the null hypothesis that there is no manipulation at the cutoff cannot be rejected at any reasonable level of significance. Frandsen (2013) has pointed out that McCrary's (2008) density test might be inconsistent for discrete running variables, such as the day count in this application. He has developed a test with preferable finite sample properties for such a scenario. Frandsen's (2013) test yields a p-value of 0.224 , thus confirming that manipulation at the cutoff is unlikely. By and large, there is no indication of any bunching or any increase in frequency just to the right of the cutoff.

A related threat to the internal validity of the design is differential sorting of offenders to either side of the cutoff. For instance, more experienced drivers could be more aware of the 365-day rule and thus sort to the right of the cutoff. Again, the data give little indication that this is the case. Figure 5a plots the average age for 3-day bins against the running variable. Age should be a reasonable proxy for driving experience. There is no sign of either a jump or drop at the cutoff or anywhere else in the distribution. Figure 5b repeats this exercise treating the percentage of female drivers in each 3-day bin as the outcome. Women generally commit substantially fewer traffic transgressions. Yet again there is no jump in the percentage of female offenders at the 366-day cutoff. This suggests that female drivers are as unaware of the 365-day rule as their male peers.⁷ Finally, I evaluate the number of prior offenses in the same way. The idea here is that habitual offenders might be more knowledgeable of the penalty catalogue and the 365-day provision and might therefore be

⁷As noted earlier, the lowess lines are merely included to show general patterns in the data. The small gap on both sides of the cutoffs is probably due to boundary bias, but there is certainly no discontinuity in the position of the point cloud.

more likely to “play” the law. But again, there is little indication of a jump or drop in the average number of priors, neither around the cutoff nor anywhere else.

Overall, there is no indication of any sorting behavior. This suggests that the 365-day rule is obscure enough to lead to a random separation of offenders into a treatment and control group, yet it is enforced to such a degree that the take-up among those who are assigned to treatment is substantially higher than among those not assigned to treatment. The treatment here is, of course, a 1-month driver’s license suspension. This gives rise to a fuzzy regression discontinuity (RD) design which is implemented as a 2SLS instrumental variable regression. One of the main challenges in this kind of setup is to pick the correct functional form. This is an important task in all empirical analyses but even more crucial in a RD setting. We want to be sure that what at first glance certainly looks like a jump in recidivism rates at the 366-day cutoff is not just a non-linearity in the data. A visual inspection of a graph that plots the outcome of interest against the running variable (such as Figure 3) provides a useful guide for picking the correct functional form. The graph reveals no obvious non-linearities. Not least for efficiency reasons, a linear functional form therefore seems to be appropriate. Nonetheless, I will also consider specifications using second and third order polynomials of the running variable and interactions of these polynomials with the treatment dummy. Gelman and Imbens (2014) show that polynomials of even higher order do more harm than good and even the cubic version of the model might be too much. It is still useful as a robustness check.

An alternative is a nonparametric approach, e.g. local linear regression (Hahn et al., 2001). Lee and Card (2008), however, argue that with a discrete running variable, such a nonparametric approach is not advisable. My running variable, the number of days between the original and the treatment transgression, is discrete but takes on many distinct values which should mitigate concerns about the nonparametric approach. Nonetheless, this method is only used as a robustness check. Section 6 will show that it leads to results that are strikingly similar to those of the least flexible parametric specification. Lee and Lemieux

(2010) also suggest that the standard errors should be clustered on the distinct values of a discrete running variable which is done throughout the paper. The second stage regression in my 2SLS model is modeled as follows:

$$Rec_i = \beta_0 + \beta_1 X_i + (\beta_2 X_i^2 + \beta_3 X_i^3) + \gamma_0 \hat{D}_i + \gamma_1 \hat{D}_i X_i + (\gamma_2 \hat{D}_i X_i^2 + \gamma_3 \hat{D}_i X_i^3) + \epsilon_i \quad (1)$$

where X_i is the running variable, that is the number of days that have passed between the date on which the penalty for the original transgression obtained legal force and the date of the treatment transgression of offender i . As is best practice in an RD setting (Lee, 2008), the running variable is centered around the cutoff. D_i is a dummy indicating whether, due to the treatment transgression, offender i had her license suspended for 1 month. This dummy is instrumented for (see below) thus the hat-superscript in equation (1). Rec_i is a dummy equal to one if offender i recidivates and commits a (third) offense within 12 months after the treatment transgression.⁸ γ_0 is the coefficient of interest and yields the treatment effect of a 1-month license suspension on the probability of recidivating within a year. In order to assess more flexible functional forms, the polynomials and interaction terms in parentheses can be added to the model. Since assignment to treatment is fuzzy, a first stage regression, yielding the predicted values $\hat{D}_i$, is necessary:

$$D_i = \delta_0 + \delta_1 X_i + (\delta_2 X_i^2 + \delta_3 X_i^3) + \pi_0 T_i + \pi_1 T_i X_i + (\pi_2 T_i X_i^2 + \pi_3 T_i X_i^3) + \eta_i \quad (2)$$

where $T_i = 1(X_i < 366)$. In other words, T_i indicates assignment to treatment and D_i indicates whether the treatment was in fact taken up. Of course, in the specifications using higher order polynomials, the first stage is constructed such that the model is exactly identified and $\{ D_i, D_i X_i, D_i X_i^2, D_i X_i^3 \}$ are instrumented for by $\{ T_i, T_i X_i, T_i X_i^2, T_i X_i^3 \}$.

All models are also run with a vector of covariates included in the regression. Controls are offender i 's age, sex, her number of prior offenses, and regional dummies for her place of residence. This provides an additional check on the internal validity of my estimates.

⁸As mentioned above, the evaluated recidivism period is extended to 13 months for those who actually receive the treatment in order to account for incapacitation effects as drivers without a license naturally have less of an opportunity to reoffend.

The covariates are balanced across treatment and control group so that the point estimates should not be affected by the inclusion of control variables. We will see in the next section that this is indeed the case.

5 Results

5.1 Reduced Form Results

The reduced form results of Table 3 yield the intent-to-treat (ITT) effect, i.e. the effect of assignment to treatment. Offenders who committed their treatment transgression within 365 days after the penalty for the original transgression had obtained legal force, are about three percentage points less likely to recidivate within 12 months than offenders who do not fall under the 365-day rule. This finding is robust to the inclusion of covariates. Changes to the functional form also have no effect on the point estimates, but result in small losses of precision. The average rate of recidivism is 25.4 percent, so these coefficients translate into a decrease in the rate of recidivism of about 12 percent. This result is also consistent with a visual analysis of Figure 3a which illustrates the ITT and would suggest an effect size of a similar magnitude.

Panel B of Table 3 shows the reduced form results with respect to speeding-specific recidivism. One might expect that future speeding offenses are deterred in particular since temporary license suspensions are imposed as a result of past speeding transgressions. Column (1) indicates that falling under the 365-day rule reduces the rate of speeding recidivism by about 2 percentage points. This translates into a reduction in speeding-specific recidivism of about 11 percent, indicating that speeding-specific recidivism is affected by the penalty just as much as any other kind of recidivism. This result is robust to including a set of covariate variables. The result is also consistent with Figure 3b which would suggest an effect of similar magnitude. The point estimate is fairly robust to changes in the functional form. However, in specifications in which a quadratic or cubic version of the running variable and the corresponding interactions with T_i are included, the standard errors increase

substantially and the coefficient is no longer statistically significantly different from zero. This is most likely a power issue. I have calculated the minimum detectable effect (MDE) using the following formula provided by Jacob et al. (2012):

$$MDE = 2.8 \sqrt{\frac{(1 - R_Y^2) \sigma_Y^2}{NP(1 - P)(1 - R_T^2)}} \quad (3)$$

where R_Y^2 and R_T^2 are obtained from a regression of Rec_i and T_i , respectively on the running variable and other covariates. σ_Y^2 is the variance in outcome in the control group, N is the sample size, and P is the proportion of sample members in the treatment group. The MDE is the smallest treatment effect that my setup has a reasonable chance of detecting. For speeding recidivism, it is 0.017 which is very close to the point estimates. As a result, we may have trouble detecting a statistically significant effect on speeding recidivism using specifications with higher order polynomials which are slightly less efficient than the linear setup.⁹

5.2 2SLS Second Stage Results

Instrumental variable estimates for the effect of a 1-month license suspension on the probability of recidivating within 12 months are reported in Table 4. The linear model without any controls suggests that a 1-month license suspension reduces the probability of committing a major traffic transgression within the next year by 5 percentage points. Adding covariate controls does not alter this point estimate substantially. A model containing an additional quadratic term of the running variable and its interaction with the suspension indicator comes to virtually the same result. The coefficient is -.052 with a standard error of .021, and it is also robust to the inclusion of covariates. The cubic model yields similar results with a coefficient of -.057 and a standard error of .024. Given a mean recidivism rate of 25.4 percent, these coefficients translate into reductions of recidivating behavior by 19 to 22 percent.

⁹The MDE for recidivism in general is 0.018, so power issues should not be a concern for the main outcome.

Columns (3) through (10) assess the robustness of my results to picking an ever smaller time window around the 365-day cutoff. For instance, column (3) focuses on offenders who committed their treatment transgression between 276 and 455 days after the penalties for their original transgressions had obtained legal force. In essence, this creates a 90-day window to both sides of the 365-day cutoff. It is comforting to see that the point estimates remain very stable. If we further zoom in on the cutoff, the point estimates continue to hover around $-.05$. Not surprisingly, the standard errors inflate substantially as would be expected since the sample size shrinks with an ever closer window around the cutoff. Wooldridge (2009) also points out that there are more than just efficiency costs to limiting the sample to observations just around the cutoff. His simulations show that this might substantially bias the coefficient of interest. Therefore, the results using a smaller time window merely serve as a robustness check. The specification that uses the full sample (i.e. column (2)) is the preferred specification.

The selection of a preferred functional form is another challenge in my setting. A visual inspection of Figure 3 finds little indication of a non-linear relationship between the running variable and the outcome. For efficiency reasons, the linear specification might therefore be preferred. I also conducted a Chi-Squared (Wald) version of the formal F-test suggested by Lee and Lemieux (2010). For this test, a set of dummies for 3-day bins of the running variables is added to the 2SLS regression. It is then tested whether this set of dummies is jointly significant. If it is not, this is a sign that the original regression (without the bin dummies) is not underspecified. If the dummies are jointly significant, then higher-order polynomials of the running variable should be added to better fit the data. The p -values of this test are reported below each coefficient and standard error pair in Table 4. The set of bin-dummies is insignificant in all specifications, indicating that higher order polynomials are not needed here. It is nonetheless comforting that both the quadratic and the cubic specification yield point estimates that are very similar to those of the linear specification. If we discount the findings from columns (9) and (10) which will likely suffer from both consistency and efficiency issues, the coefficient range stretches from -0.029 to -0.071 and is

thus hovering around the preferred -.048 estimate yielded by the linear specification using covariates and the full discontinuity sample.

The second stage results are also strikingly different from the “naive” OLS regression results of Table 6. The estimates in this table are based on the population of *all* of the about 2 million offenders who committed a major traffic violation in Germany between May 2012 and December 2012. Some of these offenses were punished with a temporary license suspension. Here, I also regress a dummy indicating whether a person commits yet another traffic violation within one year on a dummy indicating whether a person’s license was suspended following her initial transgression. Column (1) shows that there is a *positive* relationship between punishment and recidivism. That is, persons who have had their license suspended are *more* likely to recidivate. This relationship is robust to the inclusion of a set of control variables such as age and sex. Of course, nobody would argue that the “naive” OLS results document a causal relationship. Drivers who get punished might drive faster cars, be less risk-averse, or differ from unpunished drivers in many other unobservable dimensions. A simple OLS model fails to pick up these differences between punished and unpunished drivers. My regression discontinuity design, on the other hand, solves this issue because the 365-day rule ensures quasi-random assignment into treatment and control group. The treated who have their license suspended for 1 month do not differ in terms of either observable or unobservable characteristics from control members of the discontinuity sample who get to keep their licenses. In other words, the 2SLS regression yields a causal effect whereas the “naive” OLS merely reflects correlations. The difference between the two approaches turns out to be very substantial.

The results for speeding-specific recidivism are presented in Table 5. A linear model suggests that a 1-month license suspension due to the treatment transgression reduces the probability that another speeding offense is committed within 12 months by about 3 percentage points. This effect is statistically significant and robust to the inclusion of covariates. This translates into a reduction by 17 percent, given a mean rate of speeding-specific recidivism of 18.4 percent. The results for speeding-specific recidivism are, however, much

less robust to changes in specification and adjustments of the period under observation. A quadratic and cubic model using the full sample yield smaller coefficients of -0.022 and -0.020, respectively and also lead to much larger standard errors than the linear model. Moreover, if I zoom in on observations very close to the cutoff, I lose a lot of precision and some of the point estimates have a positive sign. A series of Wald-test indicates that a linear model should be the preferred model which makes the statistically significant -0.030 coefficient in column (2) the most trusted estimate. Nonetheless, their lack of robustness make the results for speeding-specific recidivism somewhat less reliable.

6 Robustness

6.1 Nonparametric Regression Results

The parametric approach that was employed in the previous section uses all observations, each of which carries the same weight regardless of how far away from the cutoff it is. However, the argument that offenders are assigned to treatment in a quasi-random fashion might be most credible for observations close to the 365 cutoff. Yet as is apparent from the previous section, limiting the analysis to observations right around the cutoff leads to both efficiency and consistency issues. Moreover, if the functional form is correctly specified, including observations that are further away does not lead to bias and improves precision.

Nonparametric techniques offer an alternative estimation approach. Instead of having the functional form determined by the analyst, a nonparametric approach ultimately lets the data pick the appropriate functional form. An example for a nonparametric technique is local linear regression (Hahn et al., 2001) which has some beneficial properties with respect to boundary bias. The idea here is to fit a weighted regression within a bin with a fixed bandwidth around the cutoff leading to a flexible fit. Observations that are close to the cutoff carry more weight than observations further away and, in fact, some observations will receive zero weight and not enter the analysis. By sidestepping the functional form issues,

local linear regression reduces the chances of bias in the treatment effect. However, it offers less precision. As mentioned earlier, nonparametric techniques should also be viewed with more scrutiny when the running variable is discrete. With discrete data, there are no observations with values of the running variable just below the cutoff (e.g. at 365.99 days in this application). Rather, clusters of observations are bunched together at uniform distances, so comparing observations “just above” and “just below” the 365-day cutoff and slowly moving across observations with slightly higher/lower values of the running variable is impossible. Nonetheless, Calonico et al. (2014) point out that in practice when a discrete running variable has a fair amount of mass near the cutoff, a nonparametric estimation might still yield a good approximation of the true treatment effect.

The key challenge with nonparametric techniques is bandwidth selection. There is a trade-off between bias and efficiency with larger bandwidths yielding more precise estimates while being more susceptible to bias. Several algorithms to pick an optimal bandwidth have been developed. They all follow a similar pattern: Starting from an initial pilot bandwidth, an iterative procedure determines an optimal bandwidth that balances the two dimensions of efficiency and unbiasedness. The two most popular algorithms by Imbens and Kalyanaraman (2011) and Calonico et al. (2014) have been implemented into software programs and are put to work in this study. Table 7 demonstrates that the nonparametric approach yields results that are strikingly similar to those of the parametric analysis. A 1-month license suspension due to the treatment transgression is associated with a reduction in recidivating behavior of 4.6 percentage points. Recall that the preferred parametric specification (column (2) in Table 4) predicted a reduction by 4.8 percentage points. The point estimate is identical regardless of the algorithm that was used for bandwidth selection. Table 7 also gives information on the number of observations in the sample that receive non-zero weights. While the algorithm by Imbens and Kalyanaraman (2011) makes use of more than 85% of observations, the algorithm by Calonico et al. (2014) uses only about half the sample. This also explains differences in precision. Figure 6a further attests to the non-sensitivity of my nonparametric results to the bandwidth choice; it shows the point estimates and 95% confi-

dence intervals over a wide range of possible bandwidth selections, including those suggested by the above mentioned procedures. The point estimates all hover around -.05 and most of them are statistically significant at the 5% level.

For speeding recidivism, the nonparametric results also resemble the preferred parametric results. In fact, the point estimate in column (4) of Table 7 is identical to the one in column (1) of Table 5. Of course, the nonparametric estimates are less precise as they are based on fewer observations. Nevertheless, the similarity in results lends additional credibility to the preferred parametric specification which predicts a reduction in speeding-recidivism by about 3 percentage points. Furthermore, Figure 6b shows that the nonparametric estimates are relatively robust to the bandwidth choice with the point estimates all falling within the -0.03 range.

Even though they should be taken with a grain of salt since, due to the discreteness of the running variable, it is very comforting that, by and large, the nonparametric results yield results very similar to those of the preferred parametric regression specification.

6.2 Recidivism Windows

In this section, I look at a variety of recidivism time windows. For one, this provides an additional robustness check and tests whether my results are sensitive to small changes in the recidivism time period under consideration. Second, it allows me to assess whether a temporary license suspension deters offenders permanently from committing further transgressions or whether over time they revert to their old behavior. It should be noted that the data put some restrictions on the length of recidivism time windows that can be evaluated. Section 3 and Figure 1 have explained the data selection process in detail. The 2-year expungement period and the lag between transgression date and data entry date put a bind on this type of analysis. The maximum recidivism period that the data allows me to reasonably evaluate is 15 months. In this analysis, offenders who committed their treatment transgression between August and November 2012 are dropped from the sample. They may

have recidivated within 15 months but their transgression might not show up in the data due to the delay in reporting.¹⁰ Obviously such a focus on longer recidivism windows also leads to substantial reductions in the sample size. On the other hand, evaluating smaller time windows allows me to use additional observations thus adding precision. It should be noted that the treated group is not forced to turn their licenses in right away. Rather, they have to deposit their license for one month at the local traffic authority within 4 months of the date on which the punishment takes legal effect. Accounting for an additional month that may be required to remind offenders that they are due to turn in their license, the smallest time window available for evaluation is thus 6 months.

A simple comparison of point estimates for different recidivism windows would be flawed since the average rate of recidivism fluctuates accordingly. For instance, the average 6-month recidivism rate is about 16 percent whereas the average 15-month recidivism rate is about 33 percent. A coefficient of, say -0.08, would thus indicate a 50 percent reduction in 6-months recidivism rates but merely a roughly 25 percent reduction in 15-month recidivism. Figure 7 therefore shows means-adjusted coefficients and confidence intervals for recidivism windows ranging from six to fifteen months. That is, the coefficients and confidence intervals are divided by the corresponding average recidivism rate. All results are obtained from a set of linear parametric regressions using covariates. The effect of a temporary license suspension is notably stable over time although the effect size is slightly larger for longer time windows. As would be expected, the point estimates also become less precise the longer the time window under evaluation as the number of available observations decreases. For all time-windows, the point estimates are negative and significantly different from zero. By and large, Figure 7a indicates that a 1-month license suspension reduces both short-run and long-run rates of recidivism by about 20 percent. This is not due to a short-run incapacitation effect but rather suggests that offenders “learn their lesson” and are permanently deterred from committing traffic transgressions.

The same exercise is repeated for speeding-specific recidivism. Figure 7b presents the

¹⁰The recidivism window for offenders who lose their license for 1 month in this instance is 16 months to account for incapacitation effects.

means-adjusted point estimates along with confidence intervals. The effect range is similar; The standard errors are slightly larger than for the general recidivism outcome. The long-run effects appear to be slightly larger than short-run effects, although these differences are not statistically significant. By and large, this graph indicates that future speeding transgressions were not deterred more than other offenses. In fact, both effect sizes and the pattern of slightly, but not statistically significantly, larger long-term effects are quite similar across Figures 7a and 7b. This similarity and the general persistence of the effect also indicate that the reduction in recidivism is not driven by the fact that the treatment group - as a result of being punished under the 365-days rule - simply becomes better informed about the peculiarities of the 365-days rule.

7 Discussion and Conclusion

The impact of punishment on future criminal behavior has always been very hard to measure. Heavy penalties are usually only handed out to offenders who have committed serious crimes. Their recidivism rates are higher to begin with and the effect of punishment cannot easily be distinguished from the effect of unobservable characteristics of these offenders such as a lower risk aversion or self-control issues. This article has exploited a special provision in the German traffic law that results in a quasi-random assignment of 1-month driver’s license suspensions to some traffic offenders but not to others. Using a fuzzy discontinuity design, I find that receiving the punishment reduces the probability of committing another offense within a year by about 20 percent. In other words, temporary license suspensions have a large deterrent effect.

This, of course, is a “local” effect in two ways. First, it is local in the sense that the estimated effect is best interpreted as the effect on offenders with values of the running variable close to the 365-day cutoff. Yet, as the summary statistics in Table 2 show, there are few observable differences between offenders close and further away from the cutoff which leaves some room for generalizations. Second, a fuzzy RD design is implemented using instrumen-

tal variable analysis. As a result, I obtain a Local Average Treatment Effect (LATE), i.e. the effect for compliers. This is the group of people who, only because they fell under the 365-day rule, had their license suspended for one month but would not have had their license suspended otherwise. The 365-day rule by its very nature only applies to repeat offenders who have committed two fairly serious speeding offenses. A 1-month license suspension may affect other types of offenders, e.g. first-time offenders, in a different way. The population of habitual offenders is, however, clearly the population that is the most interesting to policy makers so that my results might very well be seen as a case of “sometimes-you-get-what-you-want” (Angrist and Pischke, 2008). After all, it is exactly this group of repeat offenders committing fairly major transgressions that is most likely to actually get penalized with temporary license suspensions. My study obtains the treatment effects for this very population. One may also speculate that the deterrent effect for less habitual offenders might be even larger, making my estimates a lower bound for the effect of license suspensions on first-time offenders and offenders committing less severe transgressions.

The results of this article also have implications for the economic theory of crime. License suspensions due to the 365-day rule have no effect on the degree of penalties for future transgressions. As a result, they should not have an effect on the cost-benefit trade-off that rational criminals face in Becker’s (1974) framework. Yet, this study provides some compelling evidence that receiving such a penalty deeply alters future behavior. My findings could be reconciled with Becker’s (1974) model if suffering a 1-month license suspension prompts offenders, who only then realize the true costs of the penalty, to revise their expected cost estimate upwards. This kind of “learning” behavior is also consistent with economic models of crime in behavioral economics. Finally, the findings of this study are of great interest to policy makers who, as of now, have very little reliable evidence regarding the effectiveness of penalties that are levied on traffic offenders. This article suggests that temporary license suspensions for traffic offenders are an effective tool in reducing major traffic violations. Incorporating this insight into future revisions of the penalty catalogue should help prevent some of the thousands of traffic fatalities that happen everyday.

References

- AIZER, A. AND J. J. DOYLE (2015): “Juvenile Incarceration, Human Capital, and Future Crime: Evidence from Randomly Assigned Judges,” *The Quarterly Journal of Economics*, 130, 759–803.
- ANGRIST, J. D., G. W. IMBENS, AND D. B. RUBIN (1996): “Identification of Causal Effects using Instrumental Variables,” *Journal of the American Statistical Association*, 91, 444–455.
- ANGRIST, J. D. AND J.-S. PISCHKE (2008): *Mostly Harmless Econometrics: An Empiricist’s Companion*, Princeton University Press.
- BECKER, G. S. (1974): “Crime and punishment: An economic approach,” in *Essays in the Economics of Crime and Punishment*, NBER, 1–54.
- CALONICO, S., M. D. CATTANEO, AND R. TITIUNIK (2014): “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 82, 2295–2326.
- CDC (2015): “National Center for Injury Prevention and Control. Web-based Injury Statistics Query and Reporting System (WISQARS),” Tech. rep., Centers for Disease Control and Prevention.
- DEANGELO, G. AND B. HANSEN (2014): “Life and Death in the Fast Lane: Police Enforcement and Traffic Fatalities,” *American Economic Journal: Economic Policy*, 6, 231–257.
- DI TELLA, R. AND E. SCHARGRODSKY (2013): “Criminal recidivism after prison and electronic monitoring,” *Journal of Political Economy*, 121, 28–73.
- DRAGO, F., R. GALBIATI, AND P. VERTOVA (2009): “The deterrent effects of prison: Evidence from a natural experiment,” *Journal of Political Economy*, 117, 257–280.
- FRANDSEN, B. R. (2013): “Party Bias in Union Representation Elections: Testing for Manipulation in the Regression Discontinuity Design When the Running Variable is Discrete,” *Unpublished manuscript, Brigham Young University*.

- GELMAN, A. AND G. IMBENS (2014): “Why high-order polynomials should not be used in regression discontinuity designs,” Tech. rep., National Bureau of Economic Research.
- GREEN, D. P. AND D. WINIK (2010): “Using Random Judge Assignments to Estimate the Effects of Incarceration and Probation on Recidivism among Drug Offenders,” *Criminology*, 48, 357–387.
- HAHN, J., P. TODD, AND W. VAN DER KLAUW (2001): “Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design,” *Econometrica*, 69, 201–209.
- HANSEN, B. (2015): “Punishment and Deterrence: Evidence from Drunk Driving,” *American Economic Review*, 105, 1581–1617.
- HELLAND, E. AND A. TABARROK (2007): “Does three strikes deter? A nonparametric estimation,” *Journal of Human Resources*, 42, 309–330.
- IMBENS, G. AND K. KALYANARAMAN (2011): “Optimal bandwidth choice for the regression discontinuity estimator,” *The Review of Economic Studies*, rdr043.
- JACOB, R. T., P. ZHU, M.-A. SOMERS, AND H. S. BLOOM (2012): *A practical guide to regression discontinuity*, MDRC.
- JOLLS, C., C. R. SUNSTEIN, AND R. THALER (1998): “A behavioral approach to law and economics,” *Stanford Law Review*, 1471–1550.
- KESSLER, D. AND S. LEVITT (1999): “Using Sentence Enhancements to Distinguish between Deterrence and Incapacitation,” *Journal of Law and Economics*, 42, 343–364.
- KLING, J. R. (2006): “Incarceration Length, Employment, and Earnings,” *American Economic Review*, 96, 863–876.
- LEE, D. S. (2008): “Randomized experiments from non-random selection in US House elections,” *Journal of Econometrics*, 142, 675–697.

- LEE, D. S. AND D. CARD (2008): “Regression discontinuity inference with specification error,” *Journal of Econometrics*, 142, 655–674.
- LEE, D. S. AND T. LEMIEUX (2010): “Regression Discontinuity Designs in Economics,” *Journal of Economic Literature*, 48, 281–355.
- LEE, D. S. AND J. MCCRARY (2009): “The deterrence effect of prison: Dynamic theory and evidence,” Tech. rep., Princeton University.
- LEVITT, S. D. (1996): “The Effect of Prison Population Size on Crime Rates: Evidence from Prison Overcrowding Litigation.” *Quarterly Journal of Economics*, 111, 319 – 351.
- LEVITT, S. D. AND T. J. MILES (2007): “Empirical study of criminal punishment,” *Handbook of law and economics*, 1, 455–495.
- MCCRARY, J. (2008): “Manipulation of the running variable in the regression discontinuity design: A density test,” *Journal of Econometrics*, 142, 698–714.
- WHO (2013): *Global Status Report on Road Safety 2013: Supporting a Decade of Action*, World Health Organization.
- WOOLDRIDGE, J. (2009): “Estimating Average Treatment Effects: Regression Discontinuity Designs,” in *BGSE/IZA Course in Microeconometrics*.

Tables and Figures

Table 1: Excerpt from the German Traffic Penalty Catalogue: Speeding Offenses

Transgression by:	Penalization	
	Highways	Built-Up Areas
≤ 10 km/h	10€	15€
11 - 15 km/h	20€	25€
16 - 20 km/h	30€	35€
21 - 25 km/h	70€, 1P	80€, 1P
26 - 30 km/h	80€, 3P	100€, 3P
31 - 40 km/h	120€, 3P	160€, 3P, 1M
41 - 50 km/h	160€, 3P, 1M	200€, 4P, 1M
51 - 60 km/h	240€, 4P, 1M	280€, 4P, 2M
61 - 70 km/h	440€, 4P, 2M	480€, 4P, 3M
≥ 71 km/h	600€, 4P, 3M	680€, 4P, 3M

Each cell contains information on the penalization for speeding offenses. There are three types of penalties: Fines as measured in Euros (€), central registry points (P), and temporary license suspensions in months (M). A person who has accumulated 18 points will have her license revoked permanently. All points are erased if a person remains without a traffic transgression for 2 years.

Table 2: Summary Statistics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Full Sample			±90 days		±60 days		±30 days		±10 days
	[186-365]	[366-545]	[276-365]	[366-455]	[306-365]	[366-425]	[336-365]	[366-395]	[356-365]	[366-375]
Time	274.46* (51.48)	451.69* (52.28)	319.54* (25.84)	409.44* (25.95)	334.80* (17.32)	395.02* (17.19)	350.42* (8.650)	380.59* (8.691)	360.43* (2.964)	370.53* (2.813)
Age	42.66 (13.12)	43.02 (12.93)	42.88 (13.15)	42.93 (12.94)	42.78 (13.15)	42.85 (13.04)	42.84 (13.01)	42.88 (13.33)	42.85 (12.91)	42.58 (13.42)
Female	0.162* (0.368)	0.148* (0.355)	0.167* (0.373)	0.144* (0.351)	0.167 (0.373)	0.146 (0.353)	0.158 (0.365)	0.148 (0.356)	0.150 (0.358)	0.155 (0.362)
Num. Priors	2.014 (1.224)	2.004 (1.183)	2.000 (1.194)	2.018 (1.207)	1.986 (1.185)	2.019 (1.214)	1.995 (1.208)	1.993 (1.223)	1.947 (1.185)	2.003 (1.225)
South	0.227 (0.419)	0.230 (0.421)	0.231 (0.422)	0.235 (0.424)	0.227 (0.419)	0.237 (0.425)	0.217 (0.413)	0.237 (0.425)	0.220 (0.415)	0.241 (0.428)
North	0.208 (0.406)	0.209 (0.406)	0.208 (0.406)	0.207 (0.405)	0.212 (0.409)	0.208 (0.406)	0.222 (0.415)	0.208 (0.406)	0.190 (0.392)	0.212 (0.409)
East	0.179 (0.383)	0.169 (0.375)	0.178 (0.383)	0.166 (0.372)	0.176 (0.381)	0.168 (0.374)	0.170 (0.376)	0.170 (0.375)	0.168 (0.374)	0.170 (0.376)
West	0.366 (0.482)	0.372 (0.483)	0.363 (0.481)	0.371 (0.483)	0.366 (0.482)	0.367 (0.482)	0.370 (0.483)	0.364 (0.481)	0.399 (0.490)	0.364 (0.481)
Foreign	0.0200 (0.140)	0.0201 (0.140)	0.0191 (0.137)	0.0201 (0.140)	0.0183* (0.134)	0.0194* (0.138)	0.0217 (0.146)	0.0209 (0.143)	0.0230 (0.150)	0.0132 (0.114)
Recidivism (Any)	0.244* (0.430)	0.265* (0.442)	0.241* (0.428)	0.268* (0.443)	0.240* (0.427)	0.267* (0.442)	0.229* (0.420)	0.265* (0.442)	0.232 (0.423)	0.255 (0.437)
Recidivism (Speeding)	0.175* (0.380)	0.193* (0.395)	0.174* (0.380)	0.195* (0.396)	0.178* (0.383)	0.194* (0.396)	0.175 (0.381)	0.195 (0.397)	0.182 (0.387)	0.182 (0.388)
Observations	16,876	14,524	8,321	7,771	5,464	5,310	2,621	2,730	869	905

Table of means, standard deviations in parentheses. Data source is the digital German traffic registry database as of 30 April 2014. Columns (1) and (2) show descriptive statistics for the full analysis sample. In columns (3) and (4) descriptive statistics are listed for observations which have values in the running variable that fall into a 90 day window on either side of the 365-day cutoff. Columns (5) to (10) show means and standard deviations for observations with running variable values that are ever closer to the cutoff; *denotes a statistically significant difference in means at the 1% level.

Table 3: Reduced Form Regression Results

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Recidivism (Any)</i>						
Below Cutoff	−0.033***	−0.031***	−0.033***	−0.031**	−0.032**	−0.032**
	(0.009)	(0.009)	(0.012)	(0.012)	(0.014)	(0.014)
<i>Panel B: Recidivism (Speeding)</i>						
Below Cutoff	−0.020***	−0.019**	−0.014	−0.013	−0.010	−0.010
	(0.008)	(0.008)	(0.011)	(0.011)	(0.012)	(0.012)
Observations	31,400	31,383	31,400	31,383	31,400	31,383
Convariates	No	Yes	No	Yes	No	Yes
Model	Linear	Linear	Quadratic	Quadratic	Cubic	Cubic

Notes: ***/**/* indicate significance at the 1%/5%/10%-level.

Heteroscedasticity robust standard errors in parentheses, clustered by distinct values of the running variable.

Each column reports coefficients and standard errors from the reduced form OLS regression. Dependent variable in Panel A is a dummy indicating whether a person commits any traffic offense within a year of her (second) treatment transgression. Dependent variable in Panel B is a dummy indicating whether a person commits another speeding offense within a year of her (second) treatment transgression. “Below Cutoff” is the main explanatory variable and is a dummy equal to one if the treatment transgression occurred within 365 days of the day on which the penalty for the original transgression had obtained legal force. Such a person is very likely to have her license suspended for 1 month.

Table 4: Treatment Effects by Specification and Time Window - Recidivism (Any)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Full Sample		±90 days		±60 days		±30 days		±10 days	
Mean Rate of Recidivism (SD)	0.254 (0.183)		.254 (0.184)		0.254 (0.187)		0.247 (0.186)		0.244 (0.183)	
Linear	−0.051*** (0.014)	−0.048*** (0.014)	−0.050*** (0.019)	−0.048** (0.019)	−0.059*** (0.021)	−0.059*** (0.021)	−0.052* (0.027)	−0.053** (0.026)	−0.078 (0.060)	−0.082 (0.054)
p-Value Wald Test:	0.9159	0.9169	0.9859	0.9785	0.9417	0.9377	0.3196	0.3725	0.8379	0.7889
Squared	−0.052** (0.021)	−0.050** (0.020)	−0.071*** (0.026)	−0.073*** (0.025)	−0.069** (0.030)	−0.071** (0.029)	−0.055 (0.052)	−0.052 (0.049)	−0.092 (0.079)	−0.084 (0.075)
p-Value Wald Test:	0.9934	0.9895	0.9996	0.9992	0.9519	0.9479	0.3301	0.3665	0.8733	0.8612
Cubic	−0.057** (0.024)	−0.057** (0.023)	−0.047 (0.036)	−0.049 (0.034)	−0.053 (0.034)	−0.050 (0.033)	−0.037 (0.085)	−0.029 (0.074)	0.204 (0.475)	0.265 (0.603)
p-Value Wald Test:	0.9962	0.9977	0.9836	1.000	0.9854	0.9879	0.7882	0.8149	0.9865	0.9998
Convariates	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Observations	31,400	31,383	16,092	16,086	10,774	10,769	5,351	5,347	1,774	1,774

Notes: * * * / * * / * indicate significance at the 1%/5%/10%-level.

Heteroscedasticity robust standard errors in parentheses, clustered by distinct values of the running variable.

Each column reports coefficients and standard errors from the second stage of a 2SLS instrumental variable regression. The outcome variable is a dummy equal to one if a person commits another major traffic transgression within 1 year after her treatment transgression. Coefficients are displayed for the main explanatory variable, a dummy indicating a 1-month license suspension following the treatment transgression. This dummy was instrumented for using a variable for whether the treatment transgression occurred within 365 days of the day on which the penalty for the original transgression had obtained legal force. In columns (1) and (2), the entire discontinuity sample enters the analysis. In columns (3) and (4) only observations with values of the running variable within a 90-day time-window on either side of the 365-day cutoff were used. Columns (5) through (10) further limit the sample. If indicated, controls for sex, age, the number of prior offenses and a set of dummies for the region of residence were included. The reported p-values correspond to Wald specification tests. Small p-values indicate that a model is not appropriate and higher order polynomials of the running variable should be added to the regression equation.

Table 5: Treatment Effects by Specification and Time Window - Speeding Recidivism

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	
	Full Sample			±90 days			±60 days			±30 days	
Mean Rate of Recidivism (SD)		0.184 (0.387)		0.185 (0.388)	0.187 (0.390)		0.186 (0.389)		0.183 (0.387)		
Linear	−0.031*** (0.012)	−0.030** (0.012)	−0.019 (0.017)	−0.019 (0.017)	−0.035* (0.018)	−0.035** (0.018)	−0.011 (0.024)	−0.011 (0.023)	−0.010 (0.041)	−0.011 (0.039)	
p-Value Wald-Test:	0.8774	0.8470	0.9799	0.9646	0.9551	0.9514	0.4606	0.5001	0.9505	0.9473	
Squared	−0.023 (0.018)	−0.022 (0.018)	−0.045** (0.022)	−0.045** (0.022)	−0.006 (0.026)	−0.006 (0.026)	0.012 (0.040)	0.012 (0.040)	0.019 (0.046)	0.019 (0.046)	
p-Value Wald-Test:	0.9919	0.9851	0.9986	0.9974	0.9722	0.9696	0.4896	0.5043	0.9652	0.9750	
Cubic	−0.019 (0.020)	−0.020 (0.019)	0.019 (0.031)	0.018 (0.031)	0.005 (0.029)	0.008 (0.030)	0.037 (0.055)	0.044 (0.049)	0.052 (0.152)	0.072 (0.174)	
p-Value Wald-Test:	0.9758	0.9756	0.9977	1.000	0.9993	0.9993	0.8224	0.9218	0.9578	0.9544	
Convariates	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	
Observations	31,400	31,383	16,092	16,086	10,774	10,769	5,351	5,347	1,774	1,774	

Notes: * * * / * * / * indicate significance at the 1%/5%/10%-level.

Heteroscedasticity robust standard errors in parentheses, clustered by distinct values of the running variable.

Each column reports coefficients and standard errors from the second stage of a 2SLS instrumental variable regression. The outcome variable is a dummy equal to one if a person commits another major speeding offense within 1 year after her treatment transgression. Coefficients are displayed for the main explanatory variable, a dummy indicating a 1-month license suspension following the treatment transgression. This dummy was instrumented for using a variable for whether the treatment transgression occurred within 365 days of the day on which the penalty for the original transgression had obtained legal force. In columns (1) and (2), the entire discontinuity sample enters the analysis. In columns (3) and (4) only observations with values of the running variable within a 90-day time-window on either side of the 365-day cutoff were used. Columns (5) through (10) further limit the sample. If indicated, controls for sex, age, the number of prior offenses and a set of dummies for the region of residence were included. The reported p-values correspond to Wald specification tests. Small p-values indicate that a model is not appropriate and higher order polynomials of the running variable should be added to the regression equation.

Table 6: Naive OLS - Effect of License Suspension on Recidivism

	(1)	(2)	(3)
Suspension	0.010*** (0.001)	0.002** (0.001)	0.003*** (0.001)
Female		-0.074*** (0.001)	-0.079*** (0.001)
Age		-0.002*** (0.000)	-0.002*** (0.000)
Observations	2,170,728	2,170,079	2,170,079
State Dummies	No	No	Yes

Notes: ***/**/* indicate significance at the 1%/5%/10%-level.
Heteroscedasticity-robust standard errors in parentheses.
Each column reports coefficients and standard errors from an OLS regression. Dependent variable is a dummy that equals one if a person recidivates within 1 year of an initial traffic transgression. The variable suspension indicates whether a person's license was temporarily suspended after the initial traffic transgression.

Table 7: Non-Parametric Regression Results

	(1)	(2)	(3)	(4)
	Recidivism (Any)		Recidivism (Speeding)	
1-Month Suspension	−0.046	−0.046**	−0.023	−0.031
	(0.029)	(0.019)	(0.032)	(0.023)
Observations Used	13,662	26,714	9,995	16,190
Selection Algorithm	CCT	IK	CCT	IK
Bandwidth	76.44	152.2	55.62	90.61

Notes: ***/**/* indicate significance at the 1%/5%/10%-level.

Heteroscedasticity robust standard errors in parentheses, clustered by distinct values of the running variable.

Each column reports coefficients and standard errors from a local linear regression discontinuity (RD) model. A triangular Kernel function was used to construct the estimator. Coefficients yield the effect of a one month license suspension following the treatment transgression on the probability of committing another violation within 1 year. Columns (1) and (2) show the results for recidivism in general, columns (3) and (4) show the result for speeding-specific recidivism.

Bandwidth was selected using algorithms developed by Calonico et al. (2014) (CCT) and Imbens and Kalyanaraman (2011) (IK), respectively.

Figure 1: Data Selection from the Central Traffic Registry

Notes: This graph illustrates which observations can be used for the analysis. Data source is the digital German central traffic registry as of 30 April 2014. The final analysis sample is comprised of offenders with at least two speeding transgression. One, the so-called “treatment transgression” must have occurred between 1 May 2012 and 1 December 2012 (dashed frame) thus allowing for a sufficient follow-up period. Another transgression, the so called “original transgression,” must have occurred prior to the treatment transgression and the penalty for this original transgression must have obtained legal force between 186 and 545 days before the date of the treatment transgression. Persons whose transgressions over time are illustrated by circles are part of the final sample. Persons whose transgressions over time are illustrated by triangles do not become part of the final sample.

Figure 2: Treatment Probability by Time

Notes: Everybody in the analysis sample has committed (at least) two speeding violations. The first of these violations is referred to as the “original transgression.” The second of these violations is referred to as the “treatment transgression.” Loosely speaking, if the treatment transgression occurs within 365 days after the original transgression, a person should have his/her license suspended for 1 month. Data are aggregated by the count of days that have passed between the date on which the penalty for the original transgression obtained legal force and the date of the treatment transgression. Each point in this plot represents the raw probability of having one’s license suspended due to the treatment transgression against the day count. The vertical red line indicates that 366 days have passed between these two transgressions.

Figure 3: Recidivism Rate by Time

(a) Rate of Recidivism (Any)

(b) Rate of Recidivism (Speeding)

Notes: Everybody in the analysis sample has committed (at least) two speeding violations. The first of these violations is referred to as the “original transgression.” The second of these violations is referred to as the “treatment transgression.” Loosely speaking, if the treatment transgression occurs within 365 days after the original transgression, a person should have his/her license suspended for 1 month. This graph plots rates of recidivism (i.e. the probability of committing yet another, third offense) within 1 year against the number of days (in 3-day bins) that have passed between the date on which the penalty for the original transgression obtained legal force and the date of the treatment transgression. The vertical red line indicates that 366 days have passed between these two dates. The horizontal lowess lines provide a flexible fit.

Figure 4: Distribution of Time until treatment transgression

Notes: Everybody in the analysis sample has committed (at least) two speeding violations. The first of these violations is referred to as the “original transgression.” The second of these violations is referred to as the “treatment transgression.” Loosely speaking, if the treatment transgression occurs within 365 days after the original transgression, a person should have his/her license suspended for 1 month. Each column in this histogram shows the frequency (in 3-day time intervals) between these two transgressions. For instance, the first bin indicates that there were 335 incidences in which 186-188 days had passed between the date on which the penalty for the original transgression obtained legal force and the date of the treatment transgression. The vertical red line indicates that 366 days have passed between these two dates.

Figure 5: Non-Outcomes by Time

(a) Average Age

(b) Percentage Female

(c) Number of Prior Offenses

Notes: Everybody in the analysis sample has committed (at least) two speeding violations. The first of these violations is referred to as the “original transgression.” The second of these violations is referred to as the “treatment transgression.” Loosely speaking, if the treatment transgression occurs within 365 days after the original transgression, a person should have his/her license suspended for 1 month. This graph plots averages of non-outcomes against the number of days (in 3-day bins) that have passed between the date on which the penalty for the original transgression obtained legal force and the date of treatment transgression. The vertical red line indicates that 366 days have passed between these two dates. The horizontal lowess lines provide a flexible fit.

Figure 6: Treatment Effect by Bandwidth

(a) Rate of Recidivism (Any)

(b) Rate of Recidivism (Speeding)

Notes: This graph displays point estimates and corresponding confidence intervals from a local linear regression discontinuity (RD) model over a range of bandwidth choices. A triangular Kernel function was used to construct the estimator. The top figure shows the results for recidivism in general. The bottom figure shows the results for speeding-specific recidivism.

Figure 7: Treatment Effect by Recidivism Time Window

(a) Recidivism (Any)

(b) Recidivism (Any)

Notes: These graphs show point estimates and confidence bands of the treatment effect divided by the average recidivism rate for different recidivism time periods. Results stem from a set of linear regressions with controls for age, sex, number of prior offenses, and region of residence.