

Learning from Law Enforcement*

Libor Dušek[†] and Christian Traxler[‡]

January 10, 2020

Abstract

This paper studies how punishment for past offenses affects future compliance behavior and isolates deterrence effects mediated by learning. Using administrative data from speed cameras that capture the full driving histories of more than a million cars over several years, we evaluate responses to punishment at the extensive (receiving a speeding ticket) and intensive (tickets with higher fines) margins. Two complementary empirical strategies — a regression discontinuity design and an event study — coherently document strong responses to receiving a ticket: the speeding rate drops by a third and re-offense rates fall by 70%. Higher fines produce only a limited additional effect. All responses occur immediately and are persistent over time, with no ‘backsliding’ towards speeding even two years after receiving a ticket. Our evidence rejects unlearning and temporary salience effects. Instead, it supports a learning model in which agents update their priors on the expected punishment in a ‘coarse’ manner. Additional results indicate that learning from law enforcement affects drivers’ behavior more broadly, including spillovers on non-ticketed drivers.

JEL Classification: K1, D80, K42.

Keywords: Learning; Deterrence; Law Enforcement; Speeding Tickets; Regression Discontinuity; Event Study.

*This paper has greatly benefitted from discussions with and suggestions from numerous seminar and conference participants, in particular David Abrams, Ben Hansen, Matthias Lang, Michael Mueller-Smith and Arnaud Philippe. We would also like to thank Isabel Anaya, Jiri Krejsa, Marcel Tkacik, Jan Vavra and Paulina Ockajova for their excellent research assistance. We greatly appreciate the constructive cooperation of the town hall officials of Ricany and are grateful for research funding from GACR (grant 17-16583J) and DFG (grant TR 1471/1-1).

[†]Charles University, Faculty of Law, and University of Economics, Prague. libor.dusek@cerge-ei.cz.

[‡]Hertie School, Berlin; CESifo; Max Planck Institute for Research on Collective Goods. traxler@hertie-school.org.

1 Introduction

How do individuals respond to being punished for a crime or an offense? There are numerous mechanisms through which punishment can shape future behavior. Imprisonment, for instance, can imply incapacitation and peer effects. If one excludes such implicit mechanisms and considers a context in which the expected punishment for future delinquencies is held constant, past punishment would be a sunk cost: it does not evoke a *general deterrence* effect, leaving the choices of rational, perfectly informed agents unaffected (Becker, 1968).

This prediction is in stark contrast to the colloquial notion of offenders ‘learning their lesson’ from punishment – an idea that can be traced back to the classical writings of Beccaria, Bentham, and, most explicitly, von Liszt (1882).¹ To express this idea in economic terms, we consider an imperfectly informed agent; after being punished, the agent might update his priors about relevant parameters of the enforcement process (e.g. the detection risk). The agent learns from law enforcement and, in turn, adjusts his behavior. A learning framework thus captures the idea of backward-looking individuals that “are responsive to the actual experience of punishment” (Chalfin and McCrary, 2017, p.6). The objective of this paper is to identify such responses to punishment – *specific deterrence* effects mediated by learning.

Learning from law enforcement might, in principle, occur in almost any domain. Isolating learning effects, however, is very challenging. Expected punishment often increases with prior convictions. Past punishment then implies a general deterrence effect. More generally, punishment typically comprises a multidimensional treatment that influences later behavior through many different channels. Imprisonment may imply incapacitation and aging (Ganong, 2012; Barbarino and Mastrobuoni, 2014), criminogenic peer effects (Bayer et al., 2009), or negative labor market consequences (Mueller-Smith, 2015; Bhuller et al., 2019). Large fines, in contrast, may involve non-trivial income effects which can alter future behavior by influencing individuals’ risk tolerance (Yitzhaki, 1974).

Our paper avoids these issues by focusing on traffic law enforcement. We exploit a large administrative data set from a system of automated speed cameras in a suburb of Prague, Czech Republic. The data cover thousands of speeding tickets and allows us to track the driving histories of more than one million cars over several years. We observe the measured speed for every single ride through multiple speed camera zones – independently of whether a car was speeding or not. This feature clearly distinguishes our ability to track behavioral responses to punishment from studies that traditionally observe only rearrests or reconvictions. In addition, there are no relevant income effects in our context and the price for reoffending is unaffected by past punishment: fines are modest (around \$40–85) and independent of past tickets. Insurance rates remain constant, too. Moreover, punishment does not induce incapacitation (driving licenses are not revoked or suspended). The set-up is therefore ideal to isolate whether and how agents learn from law enforcement.

We provide evidence from two complementary research designs. First, we implement a regression discontinuity design (RDD) which exploits two speed-level cutoffs. The first cutoff is an

¹The latter explicitly stressed the role of punishment as a means of teaching criminals a lesson that points out the boundaries between compliant conduct and crime.

enforcement threshold used by the local police forces. If a car’s speed is above this cutoff (14km/h above the speed limit), a speeding ticket is automatically issued. The enforcement cutoff thus offers variation in punishment at the extensive margin (receiving or not receiving a speeding ticket). If the speed is above a second threshold, the fine for the speeding ticket increases from \$40 to \$85. This second cutoff offers variation in punishment at the intensive margin (tickets with low or high fines). This variation allows us to study whether higher fines influence learning and induce stronger behavioral responses.

The RDD results document strong and precisely estimated responses to speeding tickets. While driving frequencies remain unchanged, the average car’s speeding rate, i.e. the fraction of rides above the speed limit, drops by a third (from 30 to 20%); reoffending, i.e. the chance of getting a further ticket, declines by 70%. These numbers reflect a pronounced shift in the speed distribution. The mass of rides above the speed limit – including in the range below the enforcement cutoff – strongly declines. Most of this mass is shifted to the range slightly below the speed limit. As a consequence, the average speed declines by 3%,² with larger changes in cars’ top speeds.

We present a learning framework to classify these findings. Drivers have priors about the expected penalty for driving at a certain speed. After receiving a ticket, they update their priors and adjust their optimal speed accordingly. Within this framework, we distinguish between ‘fine-grained’ and ‘coarse’ updating of priors. For the former, the model predicts nuanced speed adjustments. Drivers would learn about the enforcement cutoff and, eventually, start bunching at speed levels below the enforcement cutoff. Under coarse updating, in contrast, a speeding ticket raises the expected penalty for any speed above the limit. In turn, we predict a strong drop in speeding. Heaping might occur below the speed limit but not at the enforcement cutoff. The evidence rejects the former and supports the latter predictions: drivers learn about the enforcement of speed limits but not about the threshold used in the enforcement of speeding violations.

Coarse updating further implies that, in contrast to the fine-grained case, the scope for amplifying the impact of speeding tickets through higher fines is limited. In line with this prediction, we do, on average, not find a statically significant additional effect at the second cutoff (where the fine more than doubles). Only for a theory-motivated sub-sample – rides observed under driving conditions that favour speeding – do we detect significantly stronger effects of tickets with higher fines. Despite its limited scope, this additional effect seems relevant as it affects the most aggressive speeding choices. Overall, however, the evidence suggests that the intensive margin variation in fines plays only a minor role in drivers’ updating processes.

The average response to extensive and intensive margin variation in punishment are clearly consistent with coarse updating. However, heterogeneity analyses indicate an interesting exemption: reoffenders seem better characterized as fine-grained updaters. In response to the first ticket, they adjust their speed only modestly and, ultimately, end up with a second ticket (to which they respond by a further speed reduction). Apart from this group, we detect very little behavioral heterogeneity.

²A quantitatively very similar drop in mean speed is found by Ashenfelter and Greenstone (2004) and Bauernschuster and Rekers (2019) – in the former study, in response to a change in the speed limit from 65mph to 55mph, in the latter case, after publicized crackdowns. The same studies report a 7.5% decline in traffic accidents and a 35% drop in fatality rates of accidents, respectively.

We complement the RDD results with an event study analysis which makes use of the high-frequency nature of our data. Based on information on the exact timing when speeding tickets were delivered, we explore within-car variation before and after receiving a ticket. The results corroborate all our findings from the RDD. The average treatment effects on the treated obtained from the event study are almost identical to the local average treatment effects from the RDD. The event study further reveals that the drop in speeding is immediate and very persistent over time. Over two years after receiving a ticket, there is no evidence of ‘backsliding’ towards speeding. This result contradicts Glueck (1928), who speculated that the “influence of the memory of past punishment upon the individual punished” (p.459) might fade over time such that “former punishment has [...] little, if any, lasting effect” (p.462). Our evidence clearly rejects such unlearning. The persistency suggests that drivers’ responses are not a temporary salience effect, which could emerge in the context of limited attention or cognition (Gabaix, 2019). In addition, the immediacy of the effect indicates that drivers rapidly update their priors. Only for company cars do there seem to be some minor frictions in the learning process.

In additional analyses, we study how broadly (or narrowly) drivers adjust their behavior. We show that speeding rates also drop in other speed camera zones, not only in the one that triggered the ticket. Further evidence suggests that increased compliance with speed limits inside speed camera zones is not compensated by more aggressive speeding outside these monitored zones. We also identify spillover effects. Speeding tickets not only induce a slow down among treated cars themselves but also of those traveling in lines behind them. In addition to these (at least partially mechanical) ‘backward spillovers’, we also find some evidence of ‘forward spillovers’. Speeding tickets make (otherwise aggressive) cars drive more slowly; driving less pushily, in turn, results in the car traveling ahead of a treated one driving at a slower pace, too. The latter effect, however, is weaker and less robust.

Our paper relates to several strands of research. First and foremost, we contribute to the economic literature on learning. So far, little attention has been paid to learning from and about law enforcement. Following the influential work by Sah (1991), economists have documented how own experiences and social interactions influence expectations about the criminal justice system and law enforcement more broadly (Lochner, 2007).³ There have also been studies documenting that word-of-mouth learning *between* peers influences compliance decisions (Rincke and Traxler, 2011; Drago et al., 2019). Evidence on *within*-agent learning, however, is scarce.⁴ Besides our study, we are only aware of Banerjee et al. (2019), who examine how drivers learn about the location and duration of drunk driving crackdowns in India. Consistently with our results, they find evidence of quick learning and re-optimization. In their setting, this learning pattern has interesting implications for the tradeoff between narrowly focused or more dispersed enforcement

³See also Hjalmarsson (2008). The large criminology literature is summarized by Apel (2013).

⁴For the domain of tax enforcement, one typically finds strong positive correlations between past exposure to audits and perceived auditing risk (e.g. Bérgho et al., 2018). This is consistent with randomized audit interventions which typically induce long-lasting compliance responses among audited units (e.g. Kleven et al., 2011). Obviously, the strategic interaction of tax authorities and taxpayers is more complex than our set-up. In addition to the income effects mentioned above, there are numerous mechanisms beyond learning that could shape responses to audits (DeBacker et al., 2015).

policies. We provide evidence on learning in an automated enforcement setting and our findings on coarse updating offer a novel perspective on ambiguity in law enforcement (discussed below).

We also contribute to the economic analysis of deterrence (Chalfin and McCrary, 2017). General deterrence effects have been well understood theoretically since Becker (1968) and, meanwhile, well documented empirically (e.g. Drago et al., 2009; Draca et al., 2011). For specific deterrence, the situation looks different. There is neither consensus on a formal, theoretical framework (see Nagin, 2013) nor a coherent set of empirical findings. Most economic studies of specific deterrence have focused on the impact of imprisonment – either vis-a-vis alternative, less severe sanctions (e.g. Hjalmarsson, 2009; Di Tella and Schargrodsky, 2013; Bhuller et al., 2019) or in terms of longer or harsher imprisonment (e.g. Chen and Shapiro, 2007; Drago et al., 2011; Kuziemko, 2013; Mastrobuoni and Terlizzese, 2019).⁵ Despite a large number of studies that rely on credible, quasi-experimental variation, results on specific deterrence are very mixed, with evidence from similar contexts indicating positive, negative, or null effects.⁶

We differ from (and contribute to) this strand of research in several ways. First, our set-up excludes, among others, incapacitation, criminogenic, labor market, and general deterrence effects. This enables us to isolate learning-induced specific deterrence. Second, our data offer an unusual opportunity to track behavioral responses to punishment over time precisely. This allows us to document the immediacy and persistency of the effects. Moreover, while empirical studies on criminal recidivism typically observe former offenders only when rearrested or convicted, we observe legal activities (rides that comply with speed limits) as well as illegal activities (rides above the speed limit, independently of whether a ticket is triggered or not).

Third, using a regression discontinuity and event study design, we provide causal estimates that consistently document specific deterrence effects: we find strong, immediate and persistent responses to extensive margin variation in punishment. For intensive margin variation in fines, however, we only find less precisely estimated effects that are limited to favorable driving conditions. Broadly speaking, this is consistent with the notion that variation in (learning about) detection rates is more important than the level of penalty. Fourth, we present a simple learning framework which offers a coherent interpretation of these responses to extensive and intensive margin variation in punishment.⁷ Within this framework, our findings are consistent with a coarse, discontinuous updating of priors. The evidence rejects the case of fine-grained updating as well as an interpretation in terms of temporary salience responses of agents with limited attention. As further discussed below, this differentiation has interesting policy implications. Fifth, the fact that we study variation along different margins of punishment further distinguishes our work from most other studies (Hansen, 2015, is an important exception). Note that studies on prison sentences focus mainly on the intensive margin, by comparing imprisonment with other, less severe forms of punishment (or different imprisonment conditions). Exogenous variation in punishment at the extensive margin, as it is explored in our context, is rare.

⁵See Nagin et al. (2009) for a comprehensive review of this research.

⁶Mixed results on specific deterrence are also reported in the tax enforcement literature (see fn. 4).

⁷Learning in a law enforcement context is also analyzed in Sah (1991), Lochner (2001) and Banerjee et al. (2019). We differ from these studies in that we examine a continuous rather than a binary (non-)compliance decision.

Finally, our study also adds to the economic research on traffic law enforcement. Two recent papers from this field study extensive margin variation in punishment. In an RDD that is similar to ours, Hansen (2015) exploits the discontinuity in blood alcohol content that triggers sanctions for a DUI violation. He finds a 17% decline in recidivism over a four year period. Studdert et al. (2017), who also use a within-driver design, find that experiencing a sanction reduces the likelihood of future traffic law violations by 25% during 90 days. Both studies differ from ours, in that in their setting past punishment carries a general deterrent effect (related to increases in future penalties).

Hansen (2015) also examines intensive margin variation in punishment at a second cutoff, where statutory sanctions for future offenses remain constant. He detects an additional effect of the enhanced punishment of aggravated DUI.⁸ A further RDD study with intensive margin variation in punishment is Gehrsitz (2017), who also reports an additional deterrence effect if driving license suspensions (which imply partial incapacitation) are added on top of fines and demerit points. Goncalves and Mello (2017), who use an IV strategy to identify the effect of getting away with more lenient (vis-a-vis higher) fines, find an increase in reoffending and a higher accident risk.⁹ Several features distinguish our paper from all these studies: the simplicity of our set-up (which includes neither suspensions nor jail sentences, etc.), the automated enforcement system (which leaves no scope for [learning about the] discretion of police officers; see Makowsky and Stratmann 2009; Goncalves and Mello 2017) and the fact that we can observe a continuous outcome variable, independently of whether cars are offending or speeding below the enforcement cutoff. This data feature allows us to present a precise description of a learning-induced change in the distribution of (non-)compliance behavior and long-run event study analyses, which are novel to this strand of literature.

The results from this paper have several policy implications. First, our evidence shows that automated speed camera systems are highly effective at enforcing speed limits. After receiving a speeding ticket, cars persistently reduce their speed in different speed camera zones, and there is no evidence of compensatory speeding on un-monitored parts of the road. Spillover effects on other cars further imply that the tickets' impact spreads to a broader population of drivers. Together with the speed cameras' potential general deterrence effect, this contributes to an overall decline in driving speed. These findings appear relevant, given that the WHO (2018) considers effective speed management policies as a key strategy for reducing the approximately 1.35 million annual deaths in road traffic crashes. Studies which use variation in speed limits (Ashenfelter and Greenstone, 2004; van Benthem, 2015) and their enforcement (DeAngelo and Hansen, 2014; Bauernschuster and Rekers, 2019) indicate that the observed decline in speeding might not only lower accident risks by around 10% but also trigger further positive externalities in terms of reduced air and noise pollution as well as improved emission-related health outcomes. While a comprehensive welfare analysis is beyond the scope of this paper, estimates in van Benthem (2015) suggest that the social benefits from enforcing speed limits might exceed the private costs.

⁸Consistently with a specific deterrence interpretation, this additional effect emerges only in the long-run (i.e. beyond a 3-year outcome window, which could be influenced by license suspensions, revocations, or court-ordered probation periods).

⁹This suggests that the observed additional effect from high-fine tickets – which is limited in our setting to very favorable driving conditions and, thereby, the most aggressive speeding behavior – could have non-trivial implications.

Second, and more broadly, our work points to the importance of learning and information policies in mediating deterrence effects. Support for the relevance of learning-induced deterrence effects is also provided by Philippe (2019), who studies a recent reform of minimum sentencing requirements in France. He finds that only those who have the chance to learn about the exact scope of the reform respond to it. Differences in informational policies could explain why similar policy changes can have different effects. (Re)designing institutional settings to leverage learning effects and to amplify the dispersion of information might therefore constitute an important and so far under-researched dimension of optimal enforcement policies.

At the same time, our findings on the behavioral responses around the enforcement cutoff imply that withholding certain information – and thus adding or maintaining ambiguity - may also be desirable from a policy perspective. In our context, the exact speed cutoff above which speeding triggers a ticket is unknown. Coarse updating prevents drivers from finding out the exact threshold and, in turn, implies a stronger decline in speed. With coarse updating, an ambiguous enforcement cutoff can amplify behavioral responses.¹⁰ This implication potentially extends to numerous domains where minimally harmful but illegal behavior is commonly not enforced – such as petty theft, minor drug possession, public nuisances, minor tax evasion, or environmental pollution. Hence, offenders face similar ambiguity about the exact point up to which authorities ‘tolerate’ illegal behavior and where punishment starts.

The remainder of the paper is structured as follows. After describing the institutional background and our data, Section 3 introduces our theoretical framework. Sections 4 and 5 discuss our empirical strategies and present the main findings from the RDD and the event study, respectively. Additional results are discussed in Section 6.

2 Institutional Background and Data

Ricany is a residential town with 16,000 inhabitants located just outside Prague, the Czech Republic’s capital. The town experiences heavy commuter traffic and traffic safety is a major concern. Speed measurements from 2013 suggested that 30% of all cars were exceeding the speed limit.

Speed cameras and speeding tickets. In late 2013, the city council decided to set up fixed speed cameras on five commuter roads (four with a speed limit of 50km/h, one with a limit of 40km/h; all are two-lane roads). The automated cameras record the average speed of *all cars* that pass by measurement zones of several hundred meters: cameras placed at the entry and at the exit point of each zone record every car’s number plate together with a precise time stamp. Based on the travel time, the average speed inside the zone is computed. It is important to emphasize that the speed cameras record data on all cars, independently of their speed. Note further that, while the cameras are visible, there is no ‘flash’ or any warning sign that indicates the cameras’ activity (see Figure A.1).

¹⁰On the benefits of ambiguity in an enforcement context see also Lang (2017).

All recorded data are sent to the local police electronically, in daily batches. Speeding violations that classify for a ticket (see below) are verified by an officer and then passed on to the civilian town administrators who then send the tickets to the car owners.¹¹ The speed cameras were installed in the summer of 2014 and, after a testing period, the first speeding tickets were sent in October 2014.¹²

Enforcement of fines. As in many other countries, penalties are stepwise increasing in the speed (Traxler et al., 2018). Minor speeding offenses with speeds up to 20km/h above the limit are punished with a fine of 900 CZK (approx. \$39, or 3.5% of the average monthly wage). For intermediate speeding offenses with speeds of 20–40km/h above the limit, the fine increases to 1900 CZK (approx. \$82).¹³ For the remainder of the paper, we refer to these levels as *low* and *high fine*, respectively. (Major speeding offenses, with speeds of more than 40km/h above the limit, are handled according to a different procedure. Our analysis neglects these offenses, which are extremely rare in this context.)

Several institutional aspects of the enforcement system are crucial for our research design. First, the fines do not depend on past speeding offenses. Moreover, as car insurance companies do not learn about speeding tickets, insurance do not increase either. Hence, the future ‘price’ of speeding does not increase with a speeding ticket; it remains constant. Second, unlike in Hansen (2015) or Gehrsitz (2017), punishment does not include incapacitation: tickets never result in driving licenses being revoked or suspended (or any jail sentences).

Third, the speed used to determine which penalty applies derives from an adjustment procedure that serves as a concession to prevent appeals: the measured speed is rounded down to the next integer and then reduced by 3km/h. (For example, a measured speed of 73.85km/h would be adjusted to 70km/h.) Given this procedure, the cutoff for intermediate speeding is 23km/h above the limit in terms of *measured speed*. In the remainder of this paper, we will work with the precise speed measurement, i.e. before the adjustment procedure is applied.

Fourth, when the speed cameras were set up, the local police decided to send out speeding tickets only if the *measured* speed was at least 14km/h above the limit. In contrast to the cutoff for intermediate speeding, the enforcement threshold is not prescribed in any legislation. Moreover, the cutoff was never publicly communicated (except for this paper). Hence, we do not expect drivers to anticipate the enforcement cutoff.¹⁴

¹¹The process is standardized, leaving no scope for discretion by police officers (Makowsky and Stratmann, 2009).

¹²During the early phase, there were occasional gaps in measurement and one camera was only launched in November. All issues were quickly resolved and cameras started to record 24/7 highly accurate, consistent data.

¹³About 75% pay the stipulated fines right away. Similar to a plea bargaining process, the case then ends. When the ticket is not paid, the case reverts to a sort of ‘trial’ in which the actual driver has to be proven guilty or the car owner may be found liable for a violation committed with his car by an unspecified driver. Convicted drivers or car owners are then punished by a an individually assessed fine in the range of 1500–2500 CZK for minor speeding cases and 2500–5000 CZK for intermediate cases. Convicted drivers are further punished by a deduction in demerit points. A companion paper studies the enforcement process in more detail (Dusek et al., 2019).

¹⁴We verified whether citizens requested information about the level of the enforcement cutoff from the local authority under the Freedom of Information Act. Between 2014 and 2019 there was only one request that indirectly touched on this issue. This request was made in June 2018, while nearly all observations in our analysis occurred earlier. Citizens mainly used the FoI procedure to inquire about other aspects of the speeding camera systems. (45 requests concerned statistics on the number of tickets, revenue collected, or the supplier of the speed camera technology.) The information requests and the responses by the local authority are accessible online [here](#).

Data. The city of Ricany provided us with data on the full universe of 26 million rides recorded by the speed cameras from August 2014 through August 2018. For each car, we observe its exact time of entering and exiting a camera’s zone, the measured speed (to the precision of 1/1000 km/h), and an identifier for the specific zone. The data also include an anonymized (number plate-based) ID identifying each car as well as a variable capturing the region where the car is registered. Recall that the cameras record all rides, irrespective of the speed. The data therefore allow us to observe the entire driving history (i.e. rides below and above the speed limit) of each single car ever recorded by any of the five cameras.

The speed camera data were merged with administrative data on the enforcement of the speeding tickets. We observe the date when each ticket was sent and the date it was received (i.e. when the addressee signed the delivery receipt). The data further contain the amount of the fine prescribed (900 or 1900 CZK, depending on offense severity), the payment date and, in case of non-compliance, information on further enforcement steps. Table 1 presents basic summary statistics of the full sample, split between cars that never received a ticket and cars that did. The data set covers 26 million rides from over 1.3 million cars. Among these, only 48,422 cars received at least one ticket (with a total of more than 56,000 tickets). The ticketed cars drove more frequently (84 rides on average compared to 16 for non-ticketed cars).

In our analysis, we will mainly use two outcome variables: the *measured speed* and a *speeding* dummy, indicating that the measured speed exceeds the speed limit. Summary statistics for these variables are provided in the bottom panel of Table 1. The probability that a car is speeding on a single ride is 0.126 among never-ticketed cars and 0.189 among ticketed cars. Note that the definition of speeding for this purpose includes all rides that violate the traffic law, irrespective of whether they are actually ticketed (i.e. their speed exceeds the enforcement cutoff) or not (speed above the speed limit but below the cutoff). For the measured speed, which is here normalized by the relevant speed limit (50 and 40km/h, respectively), we observe that ticketed cars drive faster (on average 5.17km/h below the limit) than never-ticketed cars (6.00km/h below the limit). As a third outcome variable, we use a (re)offense indicator which captures wheather a ride had a speed above the enforcement cutoff and thus classified for a ticket. Offending is rare: only 0.3% of all rides exceed the enforcement cutoff. By construction, the offense rate is higher (1.5%) among cars that received at least one ticket.

It is worth stressing that the data allow us to track cars but not individual drivers. Tickets are mailed to the owner of the car (who may not be the driver) and we cannot distinguish between different drivers who may share the same car (e.g. family members or company employees). While our analysis at the level of cars likely captures possible spillovers from ticketed car owners to potential co-drivers, the estimates might nevertheless represent lower bounds for within-driver responses to extensive and intensive margin variation in punishment. Another limitation of the data is that, except for the region of the number plate, we have no information about the cars themselves.

3 Theoretical Framework

This section illustrates the potential impact of speeding tickets on optimal speed choices. The analysis can easily be extended to other domains where (i) the expected punishment is convex in the magnitude of the legal violation and where (ii) offenders face ambiguities about the enforcement process, in particular, about the point up to which illegal behavior is ‘tolerated’ and goes unpunished (e.g. petty theft, minor drug possession, public nuisances).

Let the net benefits from a ride in period $t \geq 0$ with speed s_t under exogenous driving condition c_t (e.g. weather or traffic situation) be given by the function $v(s_t, c_t)$.¹⁵ $v(\cdot)$, which indicates the value of time saved or mere ‘pleasure’ from driving at this speed (net of costs of fuel consumption and accident risk), is concave in speed and $\frac{\partial^2 v(s_t, c_t)}{\partial s \partial c} > 0 \forall s, c$. At time t , the driver expects that – with probability $p^t(s)$ – speeding will trigger a ticket at costs $f^t(s)$ (fines, transaction costs, etc.). We denote the *expected costs* from the ticket, the product $p^t(s)f^t(s)$, by $q^t(s)$. This continuously differentiable function $q^t(\cdot)$ is assumed to be non-decreasing and weakly convex in s : $\partial q^t / \partial s \geq 0$ and $\partial^2 q^t / \partial s^2 \geq 0 \forall t$. Driving below the speed limit \hat{s} is never expected to trigger a ticket: $q^t(s) = \frac{\partial q^t}{\partial s} = 0 \forall s < \hat{s}, t$.

Similarly to Banerjee et al. (2019), we assume that the costs of exploratory speeding exceed its benefits. Rather than studying a bandit problem (with some value of experimentation) we thus focus on a risk neutral driver’s static problem, which is

$$\max_{s_t} v(s_t, c_t) - q^t(s_t). \quad (1)$$

For a given expectation $q^t(\cdot)$, the optimal speed s_t^* is characterized by the first-order condition

$$\frac{\partial v(s_t^*, c_t)}{\partial s_t} = \frac{\partial q^t(s_t^*)}{\partial s_t}. \quad (2)$$

3.1 Learning from Speeding Tickets

At the beginning of each period $t > 0$, a driver observes whether or not a speeding ticket was delivered. This ‘feedback’ for a ride from period $\tau < t$ with speed s_τ is denoted by $T^t(s_\tau) \in \{0, 1\}$. Based on the driver’s past driving and ticketing experience, he might then update expectations:

$$q^t(s) = P(\{s_{t-1}, T^t(s_{t-1})\}, \{s_{t-2}, T^t(s_{t-2})\}, \dots, \{s_0, T^t(s_0)\}, q^{t-1}(s)), \quad (3)$$

where $P(\cdot)$ describes the updating process, q^{t-1} is the past expectation, and $\tau = 0$ is the period of the first ride. Iterating the mapping P and accounting for the fact that a ride from period τ could, in principle, result in a ticket that is delivered in any period $t > \tau$ (see Appendix B.1), one obtains

$$q^t(s) = \Pi_t \left(\left(\{s_\tau, \vec{T}(t, s_\tau)\} \right)_{\tau=0, \dots, t-1}, q^0(s) \right), \quad (4)$$

¹⁵We neglect the possibility that drivers choose conditions, e.g. by deciding when to drive.

where $q^0(s)$ is the prior belief in $t = 0$ and the vector $\vec{T}(t, s_\tau) := (T^t(s_\tau), T^{t-1}(s_\tau), \dots, T^{\tau+1}(s_\tau))$ indicates whether and in which period a ride at speed s_τ from period τ resulted in a ticket. Current expectations are thus a mapping of past experiences and the initial expectation.

Let us now discuss different ways of learning and updating that can be captured by Π_t . A benchmark is the case of zero updating. A driver might know the ‘true $q(s)$ ’ (discussed below) and his expectations would not evolve over time. In this case, receiving a speeding ticket should not have any impact on his subsequent speed choices. This is in contrast to the predictions obtained for an imperfectly informed driver. For such a driver, a ticket provides new information regarding the probability of a speeding offense being detected (at a particular location and at a given speed s_τ) and the resulting consequences. We discuss alternative updating rules, which yield different shifts in the expected costs $q^t(s)$ that imply different responses to receiving a ticket.

Consider first a driver who, after a ride with $\{s_{t-1}^*, T^t(s_{t-1}^*) = 1\}$, adjusts expectations upward: $q^t(s) = q^{t-1}(s) + \Delta(s)$, with $\Delta(s) > 0$ in some range of s ‘near’ the previous speed s_{t-1}^* . An example consistent with this rule is illustrated by the dashed curve in the top left panel of Figure 1. The illustration shows $q^t(s) > q^{t-1}(s)$ in some speed range below s_{t-1}^* . As depicted in the bottom left panel, it then follows from condition (2) that the optimal speed $s_t^*(c)$ for constant driving conditions $c_t = c_{t-1}$ would be lower than s_{t-1}^* . In fact, as $q^t(s)/\partial s$ becomes very steep right below s_{t-1}^* , the car would drive strictly below s_{t-1}^* even in better driving conditions $c_t > c_{t-1}$.

A case with a more ‘coarse’ form of updating is depicted in the top right panel of Figure 1. The dashed line shows an example where expectations shift upwards for any speed above the speed limit: $\Delta(s) > 0 \forall s > \hat{s}$ and $q^t(s)$ becomes very steep already at modest levels of speeding (see the dashed curve $\partial q^t(s)/\partial s$ in the bottom right panel of Figure 1). The driver would drive very ‘cautiously’, at most engaging in minor violations of the speed limit, even in very good driving conditions. A further implication is that the speeding responses to such coarse updating should be larger than the responses following a ‘fine grained’ way of updating q^t .

Irrespectively of whether drivers update their expectations in a more fine grained or a more coarse manner, the framework captures backward-looking agents that “are responsive to the actual experience of punishment” (Chalfin and McCrary, 2017, p.6). As we will see below, however, the different forms of updating have important implications, among others, for the ambiguity regarding the specific enforcement cutoff.¹⁶

3.2 The ‘True’ Enforcement Process

The discussion above remained agnostic about the ‘true’ probability of getting a ticket (p) and the associated costs (f). From Section 2 we know that the actual enforcement procedure implies that the risk of receiving a ticket equals zero for any speed below the enforcement cutoff. For any speed

¹⁶In addition to updating after negative feedback (a ticket), one could also consider responses to positive feedback, i.e. updating after ‘successfully’ speeding with $s_t^* > \hat{s}$ without receiving a ticket thereafter. Judged against the expected disutility $q^t(s_t^*)$, such a ride yields a positive payoff. Note, however, that any updating response after not receiving a ticket in period $t + 1$ would be partially naïve, to the extent that speeding in period t might result in a delayed ticket that arrives in a later period $\tau > t$. If such a positive reinforcement nevertheless occurs, drivers might become more ‘optimistic’ in terms of $q^t(s)$. *Cet. par.*, this would work towards an increase in the optimal speed.

above this cutoff, however, the chance of getting a ticket essentially equals one. Agents might learn about this discontinuity. In fact, if drivers apply fine grained updating rules (as illustrated in the left panels of Figure 1), they would not stop speeding after receiving a ticket. Instead, drivers would only slightly reduce their speed. In combination with varying conditions c_t , updating and re-optimization should then induce sufficient experimentation in speed. In turn, this might enable drivers to figure out the enforcement cutoff. The optimal speed (for reasonably good driving conditions) would thus converge towards the cutoff. Empirically, we should observe bunching in the speed range below the cutoff.

This prediction is in stark contrast to what follows from coarse updating rules. The more coarsely drivers update their expectations, the more strongly they reduce their speed. With a strong shift in expectations and, consequentially, optimal speed choices, there is less scope for learning the enforcement cutoff. Under sufficiently coarse updating, we should therefore not observe any heaping at the enforcement cutoff. Coarse updating could, in contrast, contribute to bunching at or below the actual speed limit, \hat{s} . From a policy perspective this means that, under coarse updating, the unresolved ambiguity of the enforcement threshold contributes to larger behavioral responses to speeding tickets.¹⁷

As discussed in Section 2, there is a second speed cutoff at which the fine increases discontinuously. The true costs from the enforcement of speeding violations are thus stepwise increasing in speed. It is straight-forward to show that, if drivers know (or learn about) this stepwise shape, we should observe bunching below this second cutoff: instead of the interior optimum characterized by condition (2), drivers might choose a corner solution (Traxler et al., 2018; Goncalves and Mello, 2018).¹⁸ This might be even more relevant as the second cutoff is – in contrast to the enforcement cutoff – stipulated by the law and thus public. If the second cutoff were unknown, however, the discussion from above regarding the scope for learning would apply accordingly: if drivers update their expectations in a very coarse manner, they might not experiment with different levels of speeding and ultimately not learn the second cutoff either.

A further point to assess concerns the way that experiencing a ticket with either a higher or a lower fine – the intensive margin variation in punishment around the second cutoff – affects drivers’ updating and behavioral responses. If drivers update in a fine grained manner, reinforcement logic would suggest that a higher fine induces stronger updating. In terms of the illustrations in the left panels of Figure 1, there is plenty of room for $\partial q^t / \partial s$ to shift further to the ‘left’ (towards \hat{s}). If this happens, a higher fine would result in a sharper drop in speed. Under more coarse updating, however, drivers respond even to a low-fine ticket with a very conservative updating (see right panels of Figure 1). The scope for inducing stronger behavioral responses through higher fines is therefore limited. Under coarse updating, it is therefore unclear whether (and by how much) the intensive margin variation in the experienced punishment amplifies the impact from speeding tickets on subsequent speed choices.

¹⁷Conceptually, one could also interpret the result from coarse updating as a self-confirming equilibrium (Fudenberg and Levine, 1993; Battigalli et al., 2015).

¹⁸This implicitly assumes that some drivers have a sufficiently strong taste for (and conditions favouring) speeding, such that they are willing to accept a low-fine ticket.

The main behavioral predictions discussed above are summarized in Table 2. One can augment these predictions regarding the timing of responses. Firstly, within our framework, updating should produce an immediate effect. In addition, the adjustment in speed choices should be persistent as long as there is no unlearning or forgetting. Such persistency contrasts to the predictions one would obtain from models of limited attention or cognition (Gabaix, 2019). If drivers ‘forget’ or the risk of speeding tickets is simply not on their mind, there would be scope for temporary effects followed by ‘unlearning’: speeding tickets could then serve as reminders that make the enforcement system salient. Receiving a ticket could trigger a temporary decline in speeding but the effect would decay over time. Models of cognitive limitations would therefore imply some ‘backsliding’.

4 Regression Discontinuity Analyses

This section introduces a regression discontinuity design (RDD) that explores two discontinuities: the enforcement cutoff yields variation in punishment at the extensive margin (receiving or not receiving a speeding ticket); the cutoff that separates minor from intermediate speeding offenses provides variation at the intensive margin (tickets with low or high fines). After discussing the design (Section 4.1) and assessing its validity (4.2), we present the results in Section 4.3. Section 5 complements the RDD, which relies on variation between cars, with an event study design that explores within variation. The latter analysis examines the precise timing and the longevity of behavioral responses to speeding tickets.

4.1 Regression Discontinuity Design

Our raw data cover repeated observations of cars over a period of up to four years. To bring these data into a cross-sectional format suitable for an RDD, we define assignment and outcome variables that apply coherently to both cars that do or do not get a speeding ticket (see Appendix B.2). First, we compute each car’s *maximum speed*, S_i , during a given *assignment period*. This period starts with the day a car i is first observed by one of the speed cameras and ends a months later. The ride with the maximum measured speed S_i , which will serve as the assignment variable, defines the ‘trigger zone’ and the ‘trigger day’: the place where and the date on which the maximum speed was recorded. If S_i is more than 14km/h above the speed limit, the car classifies for a speeding ticket ($D_i^1 = 1$). If S_i is more than 23km/h above the limit, the higher fine applies ($D_i^2 = 1$):

$$D_i^1 = \begin{cases} 0 & \text{if } S_i < 14\text{km/h} \\ 1 & \text{if } S_i \geq 14\text{km/h} \end{cases} \quad \text{and} \quad D_i^2 = \begin{cases} 0 & \text{if } S_i < 23\text{km/h} \\ 1 & \text{if } S_i \geq 23\text{km/h} \end{cases}. \quad (5)$$

As discussed below, the two cutoffs $k = \{1, 2\}$ will translate into fuzzy treatment discontinuities.

For each day covered in our data we then identify the earliest date on which tickets that were triggered on that day were sent. From this we obtain, for any given trigger day, the earliest possible treatment day. This day defines (independently of S_i) the start of an *outcome period* of f months. Based on driving behavior during this outcome period we then compute different outcome variables

(see below).¹⁹ Our main analysis below considers assignment and outcome periods of four months ($a = f = 4$). We will document that our findings are very robust when we consider any alternative combinations of periods with $3 \leq a, f \leq 6$ months. Our main analysis observes cars during their first sequence of the assignment and outcome periods, which we denote as the first ‘episode’.²⁰

We examine either individual outcomes Y_{it} (e.g. the measured speed s_{it} for each single ride of car i at time t during the outcome period) or outcomes Y_i that are collapsed at the car level (e.g. the mean, median, 75th or 90th percentile of car i ’s speed measures during the outcome period). Analogously, we consider treatment dummies T_{it}^k , for $k = \{1, 2\}$, which indicate whether car i ’s ride t is ‘treated’: around the first cutoff ($k = 1$), treatment refers to a ride that took place after receiving a speeding ticket; around the second cutoff ($k = 2$), the treatment dummy is switched on after having received a ticket with a high (rather than a low) fine. The collapsed variables T_i^k measure car i ’s share of treated rides during the outcome period.

Not every ride of every car that classifies for treatment ($D_i^k = 1$) will be treated. We will thus observe $T_i^k < 1$ for cars with S_i above the respective cutoff from (5). On the one hand, this is due to certain cars not getting any speeding tickets (e.g. police cars, ambulances and some cars with foreign number plates have $T_i^k = 0$).²¹ On the other hand, $T_i^k < 1$ reflects variation in the tickets’ sending days: some tickets might be mailed days or weeks after the first ticket (for speeding offenses from the same day) was sent. During the early phase of the outcome period, many cars’ rides will be untreated ($T_{it}^k = 0$), resulting in $T_i^k < 1$.

Accounting for the fuzzy nature of the RDD (and, for the moment, considering only collapsed outcome and treatment measures), we estimate equations of the following structure:

$$T_i^k = \delta^k D_i^k + \kappa^k(S_i) + u_i, \quad (6)$$

$$Y_i = \tau^k D_i^k + \lambda^k(S_i) + v_i, \quad (7)$$

for both cutoffs, $k = \{1, 2\}$, and D_i^k as defined in (5). $\kappa^k(\cdot)$ and $\lambda^k(\cdot)$ are functions that capture the correlation between the cars’ assignment speed S_i and the dependent variables around cutoff k . We estimate these functions non-parametrically, using local polynomials, allowing the functions to differ on either side of a cutoff.

Equations (6) and (7) correspond to the first-stage and the reduced form of an instrumental variable approach. The first coefficient of interest, δ^k , captures the discontinuity in the treatment (or, more specifically, the discontinuous increase in the share of treated rides, once S_i surpasses the respective cutoff). The coefficient τ^k measures the reduced form effect at cutoff k . From these

¹⁹Our approach is further discussed in Appendix B.2. See, in particular, Figure B.1 which illustrates the definition of the trigger day and the outcome period for a simple example with two cars. In an earlier version of this paper we adopted a more static strategy that simply defined the initial months of the sample as assignment and latter months as outcome periods. This static approach produced similar results.

²⁰As we can track cars over multiple years, however, we can construct repeated episodes of assignment and outcome periods (where the start of a new episode is given by the first ride observed six months after the start of the previous episode’s outcome period). Multiple episodes are included in the analysis of spillovers (Section 6.2).

²¹Our data contain an identifier for certain emergency vehicles (which we exclude from the analysis). Other emergency vehicles, however, are only identifiable if they qualify for a speeding ticket. To avoid selection conditional on treatment, we do not exclude these vehicles.

two coefficients we obtain the Wald estimator for the local average treatment effect (LATE) on Y_i ,

$$\beta^k = \tau^k / \delta^k. \quad (8)$$

We will estimate the models using either *car-level* observations (i.e. for the collapsed variables as in equations (6) and (7)) or *ride-level* observations (using each single Y_{it} and T_{it}^k from the outcome period). The former approach includes just one single observation per car, irrespective of a car’s number of rides during the outcome period. The estimates thus give us the (local) effects for an *average car*. The latter approach, in contrast, puts more weight on cars with more observed rides. This will yield effects for an *average ride*.

Our main analysis of the first cutoff includes all cars during their first driving episode that had (i) an assignment speed S_i above the speed limit but below the second cutoff (i.e. $0 < S_i < 23\text{km/h}$ above the limit) and (ii) at least one recorded ride during the outcome period. When we study the intensive margin variation at the second cutoff, we analogously work with a sample of cars in their first episode with an assignment speed S_i in the range $14 < S_i < 43\text{km/h}$ above the limit (i.e. above the enforcement cutoff but below the cutoff for major speeding offenses; see Section 2) and at least one ride during the outcome period. All RDD estimates are based on rides observed between the launch of the speed cameras and July 2017.²²

4.2 Validity of RDD

4.2.1 Enforcement Cutoff

Treatment. We first provide graphical evidence on the treatment discontinuity around the first cutoff. Figure 2 plots local linear fits, confidence intervals and binned averages for the treatment rate T_i^1 , the share of ‘ticketed rides’ (after receiving a ticket) during the outcome period.²³ The graph indicates a clear discontinuity in punishment at the extensive margin: at the cutoff, the share of treated rides jumps from close to zero to roughly 80%. The underlying estimates indicate a 78.7 percentage-point (pp) discontinuity in the share of treated rides (see Column (1) of Table A.3).²⁴ As discussed above, the share of treated rides for S_i above the cutoff is below one: emergency cars, for instance, are exempted from the enforcement process and some rides during the outcome period occur before a speeding ticket is delivered.²⁵ Finally, note that the treatment rate T_i^1 is basically constant above the cutoff (rather than increasing in S_i). This suggests that the enforcement

²²At this point, the local authority reduced the enforcement cutoff by 3km/h (again without communicating anything to the public). Later rides therefore occur under a different enforcement policy.

²³The figure covers cars with an assignment speed in the range from 10km/h below up to 9km/h above the enforcement cutoff. The upper bound is motivated by the fact that the second cutoff (23km/h above the limit) is 9km/h above the first one (14km/h above limit).

²⁴Throughout the paper we report bias-corrected RD estimates with robust variance estimators (Calonico et al., 2014), implemented with the `rdrobust` package (Calonico et al., 2017). Our baseline specifications use MSE-optimal bandwidths with a triangular kernel, local linear point estimators and local quadratic estimates for the bias correction. Different kernel functions and local quadratic estimations yield almost identical results. The (in)sensitivity w.r.t. to the bandwidth choice is further discussed below.

²⁵Figure 2 also indicates a very small share of treated rides for cars with an assignment speed below the cutoff. This is due to tickets that are triggered during the outcome (rather than the assignment) period. If a car with an assignment speed $S_i < 14\text{km/h}$ is found to be speeding with more than 14km/h above the limit during the outcome

authority does not systematically prioritize sending out tickets earlier for offenses with higher speeds (within the minor offense range). We will return to this observation below.

Sorting. To validate whether the treatment discontinuity offers as-good-as-random local variation, there must be no sorting of cars below the cutoff. There are several institutional features which make sorting appear implausible in our context. First, the actual enforcement threshold is not publicly known. As we pointed out above, the cutoff is not prescribed by the law but was determined by the police once the speed cameras started working (see fn. 14). Second, and even more importantly, optimizing one’s driving speed around a given cutoff is extremely difficult in this set-up. Recall that the speed is measured in zones of several hundred meters. Hence, one would have to target a precise average speed in a zone.²⁶ Optimal targeting is further complicated by the adjustment rule discussed in Section 2 and the measurement errors of cars’ speedometers. It is therefore unlikely that drivers would be able to precisely manipulate the assignment variable S_i . As we discussed in Section 3, however, there might nevertheless be scope for figuring out the cutoff if drivers update their expectations in a fine grained manner. We might therefore observe the emergence of (at least imprecise) heaping over time.

To assess this point empirically, we first explore the density in the assignment variable around the cutoff. Neither simple visualizations nor heaping tests (McCrary, 2008) provide any evidence of sorting below the cutoff (see Figure A.2).²⁷ The data do not indicate any bunching below the cutoff, not even ‘imprecise’ bunching. Moreover, there is absolutely no evidence of the emergence of heaping over time. This last statement holds for the sample of all cars (Figure A.3) as well as for ‘local cars’ (number plates from the local region and above-median driving frequency; Figure A.4). Note that these observations are consistent with the implications of coarse updating (see the predictions from Table 2).

Balance. Next, we examine whether there are any discontinuities in cars’ observable characteristics around the enforcement cutoff (see Figure A.5 and the reduced form estimates from Table A.1). The analyses detect no systematic imbalances in pre-determined variables (such as the share of cars with number plates from the local region; the cars’ driving frequency during the assignment period (before any possible treatment); the hour/weekday/month of the trigger ride or the traffic density on that occasion).

4.2.2 High-fine Cutoff

Treatment. Evidence on the treatment discontinuity around the high-fine cutoff is provided in Figure 3.²⁸ The share of ‘high-fine treated’ rides – rides after having received a ticket with a

period (but not during the assignment period), this may result in a ticket being delivered during the outcome period. In turn, we would observe $T_{jt} = 1$ for some rides t in the outcome period.

²⁶This aspect, as well as the non-public nature of the enforcement rule, render this cutoff different from those studied in Traxler et al. (2018).

²⁷Panel (a) of Figure A.2 provides weak evidence of a minor increase in the density on ‘the wrong side’ of the cutoff. This estimate, however, is economically negligible and sensitive to bandwidth choice.

²⁸Recall that the first cutoff (14km/h above the limit) is 9km/h below the second one (23km/h above the limit). This motivates the lower bound of the S_i -range covered in Figure 3.

high-fine, T_i^2 – discontinuously increases by 81.4 percentage points (Column (1) in Table A.4). The second cutoff thus provides a strong discontinuity in the exposure to tickets with higher fines.

Sorting. As for the enforcement cutoff, we do not detect any evidence of heaping at or below the high-fine cutoff (see Figure A.6) – despite the fact that the high-fine cutoff is, in principle, public information. We also tested whether bunching would emerge over time. The data do not provide any evidence for this, which is again consistent with the prediction derived for coarse updating responses (see Table 2).

Balance. Using the same empirical strategies as in Section 4.2.1, we also examined discontinuities in pre-determined, observable characteristics around the high-fine cutoff. We do not detect any systematic imbalances (see Table A.2). At the second cutoff, however, there is a potential issue related to the timing of speeding tickets. Figure 2 above showed that T_i^1 , the share of rides after receiving *any* speeding ticket, is basically constant above the cutoff. This indicates that, *within* the range of low-fine offenses, there is no differential handling of speeding tickets with different levels of S_i . The enforcement authority might nevertheless prioritize offenses in the high-fine range and send out such high-fine tickets much more quickly. In turn, this could result in a discontinuous increase in the share of ticketed rides, T_i^1 , at the high-fine cutoff. We examined this possibility both graphically (Fig. A.7) and in reduced form estimates (Column (2) in Table A.4). The analyses indicate that there is no discontinuity in T_i^1 at the second cutoff. Hence, the variation in punishment at the intensive margin – the differential exposure to high- vs low-fine tickets as captured by T_i^2 – is the only treatment variation at the second cutoff.

4.3 RDD Results

4.3.1 Punishment at the Extensive Margin (Enforcement Cutoff)

Let us first consider responses to the extensive margin variation in punishment obtained at the enforcement cutoff. In an initial analysis, we examine possible driving frequency responses. Note that the speed cameras are positioned along commuting roads that are difficult to circumnavigate. It is thus not surprising that we find no evidence of cars either reducing their driving frequency or stopping driving in response to speeding tickets (see Columns 1 and 2, Table A.5).²⁹

Next, we turn to reduced-form evidence on speeding responses. Figure 4 shows pronounced discontinuities in the cars’ speeding rates and their mean speed during the outcome period. Cars with an assignment speed marginally above the enforcement cutoff have an 8.1pp lower speeding rate and their mean speed is about 1.35km/h slower (see Columns (2) and (4), Table A.3).³⁰ Below we will see that this decline masks stronger responses at the top of the speed distribution.

²⁹In fact, we obtain a weakly significant positive estimate suggesting that cars with an assignment speed above the enforcement cutoff are slightly *more* likely to ever return during the outcome period. While this observation is consistent with anecdotes about drivers who ‘want to see’ the cameras or drive to the town hall to complain about their speeding ticket, the effect is imprecisely estimated and sensitive w.r.t. the bandwidth choice and the length of assignment and outcome periods (a and f).

³⁰Recall from fn. 24 above that we report bias-corrected RD estimates (at the car level) with robust variance estimators (Calonico et al., 2014) under MSE-optimal bandwidths.

The Wald estimates for the LATE from receiving a speeding ticket are presented in Table 3. The estimates indicate a 9.5pp drop in the speeding rate. Relative to the rate observed in the 0.5km/h bin below the cutoff, this corresponds to a 31.8% drop (see Column 1, Table 3). Column (2) further indicates that the rate of (re)offending – i.e. the share of riders during the outcome period with speeds of more than 14km/h above the limit – drops by 70.3% (from 0.7 to 0.2%). Concerning the average speed, we find a 1.46km/h (or 3.2%) drop (see Column 3). A quantitatively very similar drop in the average speed is also reported by Ashenfelter and Greenstone (2004) (impact of a 10mph reduction in the speed limit) and Bauernschuster and Rekers (2019) (response to publicized speeding crackdowns). Our data, which allow us to look beyond the mean speed, further indicate that the decline is more pronounced at the top end of the speeding distribution: when we estimate the effect on a car’s speed at the median, the 75th- or the 90th-percentile of its speed distribution, we observe an increase in both the absolute (from 1.31 to 1.77km/h) and relative effect sizes (from 2.8 to 3.4%; see Columns 4–6, Table 3).

Figure 5 visualizes the (reduced-form) effect of receiving a speeding ticket on the speed distribution. The dashed red line depicts the speed distribution during the outcome period for cars with an assignment speed S_i within a 0.5km/h bin *above* the cutoff. Recall that around 80% of the observed rides in this group are treated (see Figure 2). Comparing this distribution with the one indicated by the green line – the speed distribution for outcome period rides of cars with an assignment speed S_i within a 0.5km/h bin *below* the enforcement cutoff – we notice a clear shift in the distribution. For cars that are marginally above the enforcement cutoff, we observe fewer rides above the speed limit. The missing mass is mostly shifted towards the mode of the distribution, which is (for both groups) roughly 3km/h below the speed limit.

Based on Figure 5, we can also compute the relative change in the speed distribution. In line with the strong drop in (re-)offense rates reported above, we obtain an approximately 50% drop in the mass of rides with a speed of 14–21km/h above the limit (see Figure A.8). Consistently with a coarse updating response, however, we observe a similarly strong drop in the range of 7–14km/h above the limit. Hence, there is a stark decline in the share of rides in the range above the speed limit but below the enforcement cutoff. This latter finding is inconsistent with a very nuanced fine grained updating of expectations.

To wrap up, both the basic estimates and the graphical evidence document that speeding tickets trigger a pronounced drop in speeding and (re-)offending. Rather than a marginal transition in the speed distribution, we detect a one-third decline in the speeding rate. Consistently with the notion of coarse updating, we observe an increased bunching mass below the actual speed limit rather than bunching at the enforcement cutoff (see Section 4.2).

Robustness. To assess the sensitivity of our estimates we first consider alternative bandwidths. Figure A.9 documents that the reduced-form effects on speeding and the mean speed are remarkably stable and significantly different from zero for any bandwidth in the range between 0.5 and 8km/h. In absolute terms, we would obtain larger (but only slightly less precise) estimates for smaller bandwidths than the MSE-optimal one.

Recall from above that our sample definition is based on ad-hoc decisions regarding the length of the assignment and the outcome period (a and f ; see Section 4.1). While the length of these periods does indeed have an impact on sample size and composition (with shorter periods, we tend to observe fewer infrequent drivers), our estimates are stable across different combinations of a and f values. This point is documented in Figure A.10, which plots Wald estimates for speeding rates and mean speed for any a and f values with $3 \leq a, f \leq 6$. (The corresponding estimates with further details on the different samples are reported in Tables A.6 and A.7.) The robustness w.r.t. these two parameters foreshadows two results from below. We will see, firstly, only a modest level of heterogeneity in observables. Secondly, the event analysis will document that behavioral responses to tickets are immediate and very persistent. The latter result implies that looking at shorter (e.g. $f = 3$) or longer outcome periods ($f = 6$) matters only in terms of sample composition.

A last important point concerns the comparison of car-level estimates from above with estimates at the level of single rides. As discussed in Section 4.1, this boils down to comparing the unweighted effect of a speeding ticket on the *average car* with the effect on the *average ride*. Estimates for the latter effects are presented in the first three columns of Table 4.³¹ Compared to the results from the collapsed analysis (see Table 3), we obtain slightly smaller point estimates, in particular for the effect on average speed. As we will discuss further below, this is due to more frequently observed cars (which gain a higher weight in these estimates) responding less strongly to tickets. In terms of relative effect size, however, the estimates still indicate a 28% [61%] drop in the probability of speeding [(re-)offending] which is similar to the relative effects observed for the average car.

Heterogeneity. Table 5 presents the results from several split-sample exercises. Columns (1) and (2) compare frequent and infrequent cars (as measured by the pre-treatment driving frequency during the assignment period). Consistently with the difference between the car- vs ride-level (or ‘unweighted’ vs ‘weighted by number of rides’) estimates from above, we find stronger responses for less frequent drivers. Both in terms of reducing the speeding rate and reducing the mean speed, cars that are observed less frequently (during the assignment period) display larger absolute and relative responses to receiving a speeding ticket.

Columns (3) to (5) compare cars according to their number plate regions. Concerning the rate of speeding, we do not find very pronounced differential responses (see Panel A of Tab. 5), although non-local cars seem to have a slightly higher speeding rate as compared to cars from the local region. For the mean speed, Panel B indicates that cars from the local region reduce their speed less strongly (in absolute and in relative terms) compared to the other cars. These findings must be interpreted with caution, however, as the ‘local’ number plates cover a relatively large area beyond Ricany itself.

In a further step, we compare the effects on rides occurring under more or less favourable traffic conditions (c_t), as captured by the traffic density (measured by the time gap to the car in front). Consistent with our basic theoretical framework, Table A.8 reports larger treatment responses under ‘good’ (above median) traffic conditions: the Wald estimates show a 2.41 and 2.78km/h (5.1 and 5.3%) drop in the mean and the 90th-percentile speed, respectively. The speeding rate drops

³¹The estimates are again robust w.r.t. different a - and f -periods, see Figure A.13.

by 15pp (37%). Under bad conditions, these estimates are much smaller (a 5pp drop in speeding rate and 0.75km/h decline in mean speed; see Panel A in Table A.8). These findings must be interpreted with caution as driving conditions in the outcome period are potentially shaped by the choice of when to drive.³²

4.3.2 Punishment at the Intensive Margin (High-fine Cutoff)

We now turn to the second cutoff, which provides variation in punishment at the intensive margin. Similarly to above, we first examine whether receiving a high-fine (compared to a low-fine) speeding ticket induces any change in driving frequency. The analyses provides null effects on circumnavigation responses (see Columns 3 and 4, Table A.5). Next, we study our two main outcome variables. Figure 6 does not indicate any visible discontinuities, either in the speeding rates or in the mean speed (see the corresponding reduced-form estimates in Columns 3 and 5, Table A.4). The Wald estimates from Table 6, which are based on a much smaller number of observations (16K cars rather than the 225K cars with an assignment speed around the first cutoff), include no statistically significant estimates either. Note that the null result is consistent with our findings from above and the predictions discussed in Section 3: under coarse updating, there is little scope for higher fines to amplify the behavioral responses to speeding tickets (see Table 2).

The data suggests that the average car does not respond differently to high- and low-fine tickets. Still, all estimates are negative and some effect sizes appear relatively large. We thus explore the sensitivity of our estimates. Concerning the different bandwidth choices, the estimates turn out to be fairly robust (see Figure A.12). When we focus on shorter outcome periods, we tend to find weakly significant effects on the mean speed (but not on the speeding rate; see Figures A.13). Replicating these estimates at the level of rides does not yield higher precision: we again obtain relatively large but imprecisely estimated effects (Col. 3–4, Table 4). A similar pattern is observed in sub-sample analyses: once again, we find no statistically significant differences in responses to tickets with higher fines (see Table A.9).

Our theoretical framework suggests that, under coarse updating, intensive margin variation in penalties has limited scope to increase the impact of receiving a ticket. Stricter punishment might nevertheless amplify the effect from a ticket if a larger fine reinforced the convexity of the marginal expected penalty, $\partial q(s)/\partial s$. This case is illustrated in Figure A.14. We should observe differential behavioral responses to low- and high-fine tickets – but only for sufficiently good driving conditions c_t that favor speeding. For average or bad driving conditions, in contrast, there would be no detectable differences.

Confronting the data with these predictions, Table 7 presents car-level estimates for good and bad driving conditions, as measured by the time difference to the next car in front on entering a speed camera zone. For rides observed in relatively dense traffic, we estimate economically small and statistically insignificant effects (Table 7, Columns 4–6). In good driving conditions, however, we observe weakly significant negative effects: a ticket with higher fines further reduces the speeding rate by an additional 8pp (21%); the average and 90th-percentile speeds drop by another 1.5km/h (3%) and 2.1km/h (4%), respectively (see Panel A, Columns 1–3 of Table 7). As

³²Using day-of-week and hour-of-day indicators for rides in the outcome period, however, we do not find any evidence that speeding tickets shape the timing of rides.

compared to the basic LATEs from receiving a speeding ticket on ‘good condition’ rides (a 15pp drop in speeding, 2.4km/h drop in mean speed; see Columns 1–3, Table A.8), these are non-trivial additional effects from experiencing tickets with higher fines.

When we condition the sample on cars that are observed under both good and bad driving conditions, the impact on the speeding rate remains significant at the 5%-level; other estimates become smaller and turn insignificant (Panel B of Table 7). Hence, the estimates provide only weak support for higher fines leading to differential updating with a higher convexity in the marginal expected penalty (see Figure A.14). In addition, as pointed out above, we must interpret these results cautiously since drivers might (conditional on the high-fine treatment) select into good or bad driving conditions. Overall, the evidence suggests that the variation in fines plays a minor role in the way drivers update and respond to tickets. The effect from intensive margin variation in fines seems limited to rides observed under favourable driving conditions.

5 Event Study

The results from the RDD provide compelling evidence on speeding responses to receiving a speeding ticket. In an event study, we now exploit the time dimension of our data to examine how quickly drivers respond and how long-lasting the effects are. The within-estimates from the event study, which yield an average treatment effect on the treated (ATT), further allow us to assess the external validity of the LATE obtained from the RDD.

5.1 Design and Sample

For each car that receives a speeding ticket, we define the treatment event by the day the first ticket is received. We refer to the ride that caused the ticket as the ‘trigger observation’. Our main sample includes cars with (i) at least one ride during a 20-week window after receiving the ticket (mirroring the 4-month outcome period from the RDD) and (ii) at least one observation (beyond the trigger) during the 12-week window before the event. (Later we will consider alternative time periods.) We further focus on low-fine tickets for speeding offenses with a speed of 14–23km/h above the limit. This allows for a meaningful comparison of the event study ATT with the LATE at the enforcement cutoff (Section 4.3.1).

Figure 7 plots the two main outcome variables in the raw data. It includes observations for all ticket events that occurred between the launch of the speed cameras and July 2017,³³ which satisfy the sample conditions described above. Week zero is defined as the last week before the first ticket was received. Each circle represents the average speeding rate (Panel a) or average speed (Panel b) of rides, binned in 7-day intervals before or after the event.³⁴ The graphs indicate strong and persistent treatment responses: after receiving a ticket, speeding rates immediately drop by around 15pp and remain almost constant over the following 20 weeks. A similar pattern is observed for the average speed, which declines by more than 3km/h.

³³As in the RDD analysis, this sample restriction accounts for the change in the enforcement cutoff in July 2017.

³⁴The sample of drivers and rides may vary between the different weeks. We address this point below.

However, Figure 7 also points to a mean reversion issue. In the raw data, the pre-treatment speeding rate gradually but distinctly increases from the 6th to the 3rd week preceding the ticket. For the mean speed, this pre-trend is even more pronounced. This pattern reflects the fact that tickets are delivered with a delay of some weeks after the offense. The trigger observations – by definition, rides with a speed above the enforcement cutoff – are thus concentrated during the weeks prior to receipt of the first ticket. This explains the pronounced increase in speeding observed in the raw data. An estimation that includes these humps would then overestimate the impact of the tickets. (A formal discussion of this point is provided in Appendix B.3.)

To deal with this issue, we exclude the trigger observations from our analysis.³⁵ The effect of this exclusion is illustrated by the lines marked with triangles in Figure 7: the massive humps disappear and pre-ticket trends are modest. The sample for our main analyses then covers 626,430 rides from 16,407 cars for their first (low-fine) ticket event. We analogously define a sample for first ticket events with a high fine. Later we will also examine second tickets.

Based on these samples we use the following specification to estimate behavioral responses:

$$Y_{izt} = \sum_{w=-12}^{20} \beta_w D_{itw} + \lambda_i + \lambda_z + \lambda_{mz} + \lambda_{dz} + \lambda_{hz} + \lambda_{ez} + \gamma X_{izt} + \varepsilon_{izt}, \quad (9)$$

where Y_{izt} is the speeding outcome of car i observed in speed camera zone z at time t . Equation (9) accounts for car (λ_i) and zone (λ_z) fixed effects. In addition, we include a rich set of dummies for time-specific effects: calendar month (λ_{mz}), day of the week and schoolday/holiday (λ_{dz}) as well as hour of the day dummies (λ_{hz}). As driving patterns differ between zones, all these dummies are interacted with the zone dummies. We also include a vector of variables capturing the driving conditions for a given ride (X_{izt}). It includes, among others, a set of dummies that capture the traffic density of a single ride non-parametrically, and weather variables (temperature, precipitation, sunshine intensity, measured at a 10-minute frequency).³⁶

The key right-hand side variables in (9) are a set of dummies D_{itw} indicating in which pre- or post-event week w an observation is recorded. Week zero, the last pre-event week, is the omitted category. The parameters of interest (the β_w 's), which are identified from within-car variation in speeding choices, have the interpretation of the expected difference in the outcome in each week relative to the last week before receiving the ticket (after partialling out other factors). In the following, we will plot the β_w -estimates together with 95%-confidence intervals based on two-way clustered standard errors (by car and by zone-hour).³⁷

³⁵This is a fairly conservative approach, as the trigger observation is in principle a relevant data-point for a car's behavior prior to the ticket.

³⁶To capture the strong influence of the traffic situation, X_{izt} includes dummies for whether the car 'ahead' of car i (at time t in speed camera zone z) entered the zone less than 2, 2–4, ..., 18–20, or more than 20 seconds prior to car i . As a second measure, we also include the total number of cars passing through zone z in a particular hour of that day. The weather data were collected at the meteorological station at the Research Institute for Landscape and Ornamental Gardening, located in a small town 7km away from Ricany. These variables vary only over t but not between z , as the weather conditions are practically identical at all five speed camera locations.

³⁷Clustering only at the level of cars yields similar standard errors.

5.2 Event Study Results

5.2.1 Response to Punishment at the Extensive Margin

Panel (a) of Figure 8 plots the estimated coefficients and confidence intervals for the binary speeding outcome. The effects on the weeks prior to receiving the ticket exhibit no pre-trend. The baseline rate of speeding, that is, the average speeding rate during the last week prior to receiving the ticket, is 27% (bottom panel of Table A.11, Column 1). Immediately after receiving the ticket, the speeding rate drops by 7.4pp. It further declines in the 2nd (and, to a lesser extent, in the 3rd) week after receiving the ticket. These effects are precisely estimated, with the width of the 95% confidence intervals being less than 2pp. The decline in the speeding rate stabilizes at about 10pp below the pre-ticket level (even though there is a slight but statistically insignificant downward trend). Relative to the pre-ticket baseline, the 10pp drop implies a 37% reduction in the probability of speeding.

Panel (b) presents the analogous estimates for the measured speed. We observe a similar pattern: an immediate drop in measured speed by 1.0km/h in the first week with an additional reduction in later weeks. Over the 20 weeks, the estimated effect size varies but remains approximately flat in the range of 1.2–1.4km/h below the pre-treatment speed. Relative to the baseline average of 44.86km/h in week zero, this effect implies a reduction in speed by 2.7–3.1%.

The two figures establish the key finding from the event study design: the behavioral responses to receiving a speeding ticket are immediate and persistent over 20 weeks. In terms of absolute and relative effect sizes, it is worth noting that the ATT estimates from the event study design are very similar to the LATEs found in the RDD analysis. Note further that we do not find any evidence of ‘backsliding’: speeding outcomes do not revert towards the pre-ticket levels. This clearly rejects the idea of ‘unlearning’. The large drop in speeding rates further supports the notion of coarse updating introduced in Section 3.

Further evidence along these lines is provided in Figure 9, which depicts the effect on the speed distribution. It is analogous to Figure 5, except that it is based on a within-car comparison: the two lines compare the cars included in the event study sample before and after receiving the ticket. The speed distribution for post-treatment rides (dashed red line) contains significantly less mass in the range between the speed limit and the enforcement cutoff than the pre-ticket distribution (solid green line). That mass is mainly shifted to speeds about 5km/h below the speed limit, the mode of the pre-ticket distribution. Such a shift is inconsistent with fine-grained updating which would imply an increase in the mass below the enforcement cutoff.

Figure 8 indicates that effects are persistent over a 20-week period. To explore whether there is any backsliding in the long-run, we estimate an alternative specification that (i) extends the time window to 6 months before and 24 months after the ticket, (ii) replacing the weekly dummies D_{itw} from equation (9) with monthly dummies. With such a long horizon, compositional effects are an issue: observations far away from the ticket date (both before and after the event) would be disproportionately composed of regularly driving cars that may differ in their speeding pattern and treatment responses to tickets. We therefore (iii) restrict the sample to ‘regular’ cars that

have at least one observation in each 3-month interval during the 6 + 24 month sample window. Inevitably, the sample includes fewer cars (4,291) but a large number of rides (991,333).

The estimated coefficients on the monthly dummies are plotted in Figure 10 (and reported in Table A.12). For the speeding rate, the estimates are remarkably similar to the weekly estimates, both in terms of the qualitative pattern and the effect sizes. There is absolutely no evidence of backsliding. On the contrary, the effect size slightly increases over the two-year outcome window. This seems to reflect the general decline in speeding observed by the speed cameras, which is not absorbed by our control variables. For speed, the estimates exhibit a visible pre-trend, suggesting that these cars increase their average speed over time before eventually getting a ticket (see fn. 16). However, the drop in average speed after receiving a ticket is again similar in magnitude to the weekly estimates. Over the two-year follow-up, there is no backsliding but a further decline in speed. All in all, these long-run estimates provide no evidence of any decaying of the effects over time. Within our theoretical framework, the estimates are consistent with a permanent update of the expected costs of speeding, with no ‘unlearning’.³⁸

5.2.2 Response to Punishment at the Intensive Margin

Analogously to the RDD, we next investigate whether there is an additional effect from receiving a speeding ticket with a higher penalty. To do so, we estimate equation (9) for a sample of cars whose first ticket carried the high fine. The results – together with our estimates for low-fine tickets – are presented in Figure 11. For the speeding rate, the average effect sizes are virtually identical for cars receiving high- or low-fine tickets. For measured speed, the effects of a high-fine ticket range between 1.5–2.0km/h, which is more pronounced than the corresponding effects from a low-fine ticket (see Panel (b) of Figure 11). However, the estimates are less precise and typically overlap with those obtained for low-fine tickets. Moreover, the pre-ticket baseline speed is also slightly higher (45.75 rather than 44.86km/h) in the high-fine sample.³⁹

Recall that the observation that higher fines hardly amplify responses is consistent with our predictions for the case of coarse updating (Table 2). As discussed in the RDD analysis in Section 4.3.2, however, an intensive margin increase in punishment may result in a more convex marginal expected cost function $q(s)$. As depicted in Figure A.14, this would imply scope for differential responses to high- and low-fine tickets under good driving conditions that favour speeding. As in the RDD, the event analysis provides some support for this case: under good driving conditions, the drop in the measured speed is more than 1km/h larger for high-fine than for low-fine tickets (see Figure A.15). Concerning the probability of speeding, however, there is no difference. Hence, the additional effect from higher fines seems limited to the speed level choices under good driving conditions. Given that this affects the most aggressively speeding rides, however, this could have non-trivial implications (Goncalves and Mello, 2017).

³⁸Admittedly, the persistent salience of the enforcement regime might simply be due to the visibility (and stable functioning) of the speed cameras, which serve as constant reminders.

³⁹In addition, there appears to be an upward trend in speed during the last pre-ticket weeks in the high-fine sample, which further complicates the comparison of effect sizes.

5.2.3 Heterogeneity Analyses and Extensions

To analyze whether and how the effects vary across different types of car owners, we follow the RDD analysis and first compare frequently and infrequently observed cars (as measured by the cars' pre-ticket driving frequencies). The results, which are presented in Figure A.16, again corroborate the RDD estimates: we observe slightly larger effects for less frequent cars, both for the speeding rate and the driving speed. In the same vein, Figure A.17 compares cars according to their number plate region. Consistent with the RDD results reported in Table 5, we observe smaller effects for cars from Prague and the local region than for cars from other regions of the country. The differences, however, are in general statistically insignificant.

The event analysis, which in contrast to the RDD focuses only on cars that receive a ticket, enables us to explore further dimensions of heterogeneity. Among these, we observe whether a ticket was sent to a physical person or a 'corporation'.⁴⁰ If the car owner is a private person, the individual (who is also in charge of paying the fine) learns about the speeding ticket directly. For cars owned by corporations, in contrast, there might be more frictions in the learning process. Corporations might pay the fines on behalf of their drivers without informing them. In cases where multiple drivers share one car, it might be hard to identify the responsible driver. Even if the message reaches the (relevant) driver, it is unclear whether all the information included in the speeding ticket (e.g. regarding the specific location) is accurately communicated. We therefore expect private cars to respond more swiftly and more strongly to speeding tickets than drivers of cars owned by corporations. Figure 12 provides some support for this expectation. During the first 3–5 weeks after receiving the ticket, there is a 2pp stronger drop in the speeding rate and a 0.5km/h larger decline in speed (with some differences being statistically significant; see Table A.13). In later weeks, however, the differences shrink. This 'catching up' could be explained by a delayed communication and information transmission process, which might be particularly relevant for larger corporations.

A final dimension of heterogeneity concerns whether the car owner did or did not pay the speeding ticket (within 90 days).⁴¹ Obviously, this is an endogenous rather than a pre-determined characteristic. We thus have to be cautious in interpreting the strong heterogeneity in ticket responses documented in Figure 13. The estimates indicate that cars who pay their tickets slow down much more strongly. Those who do not pay their tickets nevertheless do also adjust their driving behavior. During the first three weeks, the speeding rate among the former group drops by 8–11pp. Among the latter, the drop amounts to a mere 3–5pp. Over time, this gap narrows but it does not fully disappear in later weeks. A similar (but less precisely estimated) pattern is also observed for speed level.

Sensitivity Analyses. In a first set of robustness checks, we modified equation 9 by excluding/including alternative measures of traffic and weather conditions, by excluding observations

⁴⁰The term 'corporation' is used as a shortcut that encompasses all judicial persons, e.g. business corporations, partnerships, non-profits, and various public entities. Note, however, that a non-negligible fraction of single-person businesses are legally organized as limited liability partnerships with only one partner, and in such cases, the judicial person *de facto* represents an individual.

⁴¹15% of car owners do not pay their tickets within that time. Conditional on the payment being made, the average time from receiving a ticket to paying the fine is 10 days.

during highly congested traffic conditions, and by using alternative ways of controlling for long-term trends in speeding (linear and polynomial trends, month fixed effects). These alternative specifications produced effect sizes virtually identical to our main specification.

Our second set of robustness checks takes a very different approach. The baseline regression (equation 9) implicitly models behavioral response as a function of time. It may be the case that the underlying learning mechanism is associated with the actual engagement in the activity (driving through camera zones). That is, the effects might kick in as agents make speeding choices for the 1st, 2nd, etc. time after being punished. To account for this, we replicate the event study with treatment dummies defined by the order of rides. We sort rides for each car and then define dummies grouped over intervals of five rides before and after the car’s ticket. In a specification akin to equation (9), we then include dummies that cover the sequence up to the 100th ride occurring after the car received a ticket, and up to 70th ride before the ticket. The results of this exercise are reported in Figure 14, where the omitted category is now given by the last five rides before the ticket.

We again observe a large and persistent negative impact on both speeding measures. Quantitatively the estimates are very similar to our basic results obtained with weekly dummies. A noteworthy difference is the clear positive pre-trend, which suggests that drivers explored higher and higher speed levels until they received a ticket.⁴² In addition, we observe a modest increase in the effect size between the 5th and the 20th post-ticket rides. Both observations, however, are partially shaped by changes in sample composition: observations further away from the ticket are increasingly composed of cars with higher driving frequency.

Second Ticket. So far, we have focused only on the impact of the first speeding ticket. Among the 33,016 cars that received at least one ticket, however, 17.52% also received a second ticket. The probability of reoffending is significantly higher for cars owned by corporations. There is also a strong, positive correlation with the measured speed of the first ticket, suggesting that cars with a stronger taste for speeding select into this small group (6K out of a total of 1.3 million cars).

To further examine how reoffenders differ from other groups, we present event study estimates that compare cars that did or did not reoffend after their first ticket. While one has to keep in mind that this sample split is based on an outcome variable, the patterns depicted in Figure 15 are interesting and provide new insights into the heterogeneity in the forms of updating. While the responses of the cars that never reoffend are basically identical to our main results for the average car, the reoffenders seem to adjust much more gradually: during the first 10 weeks after their (first) ticket, their speeding rate drops by only 3–6pp and their mean speed declines by less than 0.5km/h. One interpretation of these smaller responses is that the cars that later reoffend initially update their expected costs in a much more fine-grained manner. These cars are more likely to continue speeding and so, eventually, end up with a second ticket. Hence, while the average response to speeding tickets is shaped by a majority of coarse updating drivers, there appears to be a small group whose responses are consistent with more fine-grained updating.

⁴²This observation is consistent with updating to ‘positive feedback’, as discussed in fn. 16.

The behavioral adjustments of reoffenders, depicted in Figure 15, can be partially driven by the arrival of the second ticket.⁴³ We examine the impact of the second ticket on reoffenders separately in Figure 16. While the sample is highly selected and much smaller, the baseline pre-ticket mean outcomes are not too different from our main sample.⁴⁴ Compared to their response to the first ticket (Figure 15), the reoffenders adjust immediately and more strongly to the second ticket: we observe, on average, a 7pp drop in the speeding rate and a 1km/h drop in the mean speed. Relative to the week zero baseline, this corresponds to a 25% (2%) decline in speeding (speed), which is clearly below the ATT effects estimated for the first ticket. These findings suggest that reoffenders, who only modestly adjust their speed after the first ticket, seem to learn their lesson after the second ticket. Even after the second ticket, however, the group continues to have a 20% speeding rate, which is well above the post-treatment speeding rate of 16% observed on average. This observation is again consistent with fine-grained updating which implies – in contrast to more coarse updating – a higher frequency of optimal speeding choices above the speed limit but below the enforcement cutoff.

6 Further Results

6.1 Narrow or Broad Learning?

The results presented above coherently document drivers’ responses to facing punishment. Our evidence rejects the case of no-updating and supports the learning and (coarse) updating framework described in Section 3. This subsection now explores whether drivers’ learning is more ‘broad’ or ‘narrow’. More specifically, we ask whether drivers solely update $q_z^t(s)$ – the expected costs of speeding in the camera zone z that triggered the ticket – or whether they update their expectations $q_\ell^t(s)$ at other locations $\ell \neq z$, too. We present two empirical strategies to address this question.

Our first approach compares the impact of a ticket triggered in zone z on rides observed in the *same* and in *other* zones. Figure 17 plots event study estimates for this comparison.⁴⁵ We observe behavioral responses both in the same zone and also in other zones, with the drop in the speeding rate and in measured speed being significantly larger in the zone that triggered the ticket. This pattern is replicated in the RDD estimates, which are reported in Table 8. It is important to note, however, that the baseline rate of speeding (roughly 40% vs 19%) and the baseline level of speed (47.6 vs 44.2km/h) are larger in the trigger zone as compared to the other zones. This is intuitive, as cameras on faster roads generate (*cet.par.*) more tickets. If we account for this fact by computing relative effect sizes, we observe much more similar effects. The RDD estimates, for instance, indicate a 28.9% decline in speeding in the same zone and a 31.9% drop in other zones (see Panel A, Columns 1–2, Table 8). The results are similar when we constrain the sample to cars that are observed in both the same and other zones (Columns 3–4). For the relative declines in

⁴³28% of these cars receive their second ticket within the 20 weeks period.

⁴⁴The 1,694 cars included in Figure 16 are a subset of the 2,551 reoffending cars covered by Figure 15. The difference is explained by the sample restriction, which focuses on cars with at least one non-trigger observation in the 12 weeks before and at least one observation in the 20 weeks after the second ticket.

⁴⁵Note that the camera that triggered the first ticket varies between cars.

the measured speed the initial gap remains. The RDD estimates indicate a 4.0% drop in the same and a more modest 2.4% decline in other zones (see Panel B, Columns 1–2, Table 8; Columns 3–4 report similar effects for a constrained sample of cars). The relative effects implied by the event study estimates are similar (see Table A.11).

These findings document that drivers seem to learn more ‘broadly’. They adjust their behavior not only at the place of the past offense – where they incurred a law enforcement response – but also become more compliant on other roads monitored by speed cameras. A natural follow-up question is whether these cars would also adjust their speed on roads that are *not* monitored.

To tackle this question, we exploit the fact that our data contain the exact time when each car exits from (the endpoint of) one speed camera zone and enters into (the start of) another one further down the road. Based on this time gap, we can learn about the cars’ speed in the unmonitored stretch in between: the faster a car drives, the sooner it enters the second camera zone. In principle, one can derive three different predictions about the impact of a speeding ticket: a first hypothesis is that, while cars update $q_z^t(s)$ in different zones z (see above), they would not do so on roads not monitored by speed cameras. In this case, we should not see any change in the time spent on the unmonitored parts of the roads. A second hypothesis is that the learning spills over and results in a broad adjustment in expected fines, beyond roads covered by cameras. One would thus expect a drop in the inter-zone travel times. (In part, this might also happen mechanically, as drivers exit the first zone at a lower speed.) Finally, a third hypothesis is that drivers’ optimal speed choices are influenced by travel time targets. After slowing down (i.e. losing time) in the speed camera zone, they might want to ‘catch up’ by driving faster on the unmonitored part of the road.⁴⁶ We should then observe faster inter-zone rides in response to a speeding ticket.

Among the five speed camera zones, there is only one combination where leaving one zone leads (after a left turn) into the entry point of another one. This ‘unmonitored trip’, however, is rarely observed. Moreover, drivers encounter a traffic light on the way, which introduces sizable variation in the travel time on the 1,080 meters between exiting the first and entering the next zone. We can nevertheless use the measured time to compute the average speed on the trip between the zones. Focusing on the variation around the enforcement cutoff, we then apply the RDD strategy from Section 4 and estimate reduced-form effects on these inter-zone rides.

The estimates clearly reject the third, ‘catch-up’ hypothesis: car-level estimates yield a negative but very imprecisely estimated effect on the mean speed;⁴⁷ for the 90th percentile speed, we obtain a weakly significant negative coefficient (see Columns (1) and (2), Table A.10). In ride-level estimates, the effect on speed is much smaller and insignificant (Column 3). When we consider the log travel time as a dependent variable, the effect is once again significant and the positive sign indicates that cars slow down. All these estimates, however, are fairly sensitive (e.g. regarding bandwidth size).

To wrap up, our analyses provide evidence that is consistent with the first (and weakly supportive of the second) hypothesis set out above. We do not observe any ‘catching-up’ attempts

⁴⁶In terms of our model, the marginal benefit of speeding might increase after having slowed down before.

⁴⁷The average speed for this trip is far below the speed limit, presumably due to the traffic light.

on the unmonitored part of the road; there seems to be either no response or a minor decline in the drivers' speed. Together with these results from the between-zone comparison, our evidence weakly supports the notion of broad learning and behavioral responses to law enforcement.

6.2 Treatment Spillovers

Our last step is to analyze the potential spillovers of speeding tickets. The basic idea is straightforward: if a 'ticketed' car slows down, the following car might slow down, too.⁴⁸ In fact, in dense traffic conditions, such spillovers might reach beyond the next car in line. In addition, we also explore spillovers on the car ahead. Given that a ticket makes an (otherwise aggressive) car drive more slowly, its less pushy driving might also affect the car in front of the 'ticketed' car.

To evaluate such spillovers we identify *groups* of cars (g) in our data. In particular, we consider lines of two or more cars which all enter a given camera zone within 5 seconds of each other. The first car (the 'front' of the line) is one that enters the measurement zone at least 10 seconds after the car ahead of it. The last car of the line, which marks the 'end' of the group, enters the zone more than 5 seconds ahead of the next one. Based on these definitions we can then zoom into the sequences with different lines.⁴⁹ We study the responses of cars (rides) in position j in group g to the treatment of the car at position ℓ in the same group g . We then examine the *spillovers* from a speeding ticket for a car, e.g. in position $\ell = 2$ on any subsequent car in the line (with position $j > 2$) but also on the car in front (in position $j = 1$). The estimates for cars with position $j = \ell$ allow us to compare the *direct* treatment effect on cars in different positions j within a line g .

To estimate these effects we augment the RDD from Section 4. We run ride-level estimates for the speed outcome $Y_{i(j)gt}$ from car i observed in position j in group g during a given ride t :

$$Y_{i(j)gt} = \tau_j^\ell D_{\ell g} + \lambda^\ell(S_{\ell g}) + v_{i(j)gt}, \quad (10)$$

The key parameter, τ_j^ℓ , measures the reduced-form effect of having a car in position ℓ that has an assignment speed above the enforcement threshold on the outcomes of (other) cars observed in position j within the same group g .⁵⁰ To obtain the Wald estimate $\beta_j^\ell = \tau_j^\ell / \delta_j^\ell$ we complement the reduced form with the corresponding first stage,

$$T_{\ell gt} = \delta_j^\ell D_{\ell g} + \kappa^\ell(S_{\ell g}) + u_{\ell gt}, \quad (11)$$

where the treatment dummy $T_{\ell gt}$ indicates whether ride t of the car in position ℓ in group g is 'ticketed' (i.e. occurs after receiving a speeding ticket). Two remarks are necessary here. First, for a given j and ℓ , β_j^ℓ is identified from between group variation in $D_{\ell g}$ (and $T_{\ell gt}$) driven by $S_{\ell g}$. For

⁴⁸We would like to thank Ben Hansen for highlighting this idea.

⁴⁹The estimates presented below are qualitatively robust to changes in these specific definitions. On a more conceptual point, note that groups (or lines) are defined around single rides. Hence, the rides at t and $t + 1$ from a given car i are, by definition, partitioned into separate groups. It should also be clear that the composition of (and sequence within) a group naturally changes over time. This also means that the time subindex t used below is redundant and serves only to illustrate (consistent with the notation from Section 4) that we focus on ride-level rather than collapsed outcomes.

⁵⁰Analogously to D_i^1 from (5), $D_{\ell g}$ equals one if $S_{\ell g} \geq 14\text{km/h}$ above the limit.

$j = \ell$, the estimates are conceptually comparable to the ride-level estimates presented in Table 4. For $j \neq \ell$, the RDD exploits variation in *other* cars' assignment speed $S_{\ell g}$ rather than the car's 'own' S_i . Second, we estimate β_j^ℓ separately for all rides observed in position $j \in \{1, \dots, 5\}$. Hence, any regression includes just one observation per group but, in general, repeated observations from a given car i (see fn. 49). We thus cluster standard errors at the car level.

The results of this RDD are presented in Tables 9 and 10.⁵¹ Panels (a) – (d) decompose the β_j^ℓ -estimates for groups of cars with two, three, four, five or more cars in a line. The different columns present effects on the $j = 1^{\text{st}}, 2^{\text{nd}}, \dots, 5^{\text{th}}$ car within such lines. Let us first discuss the case $\ell = 1$, i.e. where the first car within a line is treated (see Columns 1 – 5). Unsurprisingly, the estimates document direct treatment effects on the (treated) cars in position $j = \ell = 1$. For lines with two or three cars (Panel (a) and (b) in Tables 9 and 10), the direct effects are statistically significant and quantitatively very similar to the ride-level estimates reported in Table 4, Columns (1) and (3). In addition, the estimates also reveal spillovers on the 2nd and 3rd cars traveling in these lines. For groups of three cars (Panel b), for instance, the speeding rate of the first (i.e. the 'ticketed') car drops by 7.4pp, and that of the 2nd and 3rd cars by 5.5pp and 5.1pp, respectively. A similar pattern is observed for the driving speed.

Next, we consider groups where the car in position $\ell = 2$ is ticketed (Columns 6 – 10). In addition to the direct effect on cars in this position $j = 2$, we again find some evidence of spillovers. However, both the direct effects and those on the following cars tend to be smaller. This seems to be due to the fact that the treated car's scope for speeding is already constrained by the mere fact that this car is observed in position $\ell = 2$ within a line of relatively dense traffic.⁵² Despite that, we obtain some weak evidence of small treatment spillovers on the rides in front of the treated car (i.e. in position $j = 1$; see Column 6, Table 10). Weakly significant estimates of such 'forward spillovers' are also documented in Columns 11 – 15, for the case where the third car ($\ell = 3$) in a line is treated. As before, however, the estimates show weaker direct effects (again, in a context where the level of speed and the rate of speeding is already constrained by the traffic situation) and also smaller and less precisely estimated 'backward spillovers'.

To summarize, our estimates provide evidence of treatment spillovers. Especially for cases where the first car in a line is treated, the ticket induces a decline in the speed and speeding rate of (at least) the next two cars in the line. In addition to these (partially mechanical) backward spillovers, we find some evidence of 'forward spillovers' on the car ahead. The latter spillover

⁵¹The sample includes observations from all groups of two or more cars for which a car in position ℓ was observed (i) during the outcome period with (ii) an assignment speed $S_{\ell g}$ around the enforcement cutoff. The sample size thus varies between (but not within) different positions ℓ and for different group sizes (number of cars in a line).

⁵²For, e.g. lines with three cars where the one in position $\ell = 2$ is treated, the average speed/speeding rate of rides in position $j = 2$ (of cars with an assignment speed marginally below the cutoff) is 43.9km/h / 18.5%. For lines where the first car ($\ell = 1$) is treated, the corresponding averages among rides in position $j = 1$ amount to 45.8km/h / 31.0% (see the line Y(left) in Panel (b), Columns 1 and 7 in Tables 9 and Table 10). A further observation worth noting is that the cars that lead lines – which, by our definition of car groups, have a free road ahead – drive faster than cars following them. For instance, within lines of four cars (in which the leading one is ticketed), the speeding rate of (marginally untreated) cars in position $j = 1, \dots, 4$ monotonically declines from 28.4, 17.5, 13.6 to 10.3% (see Y(left) in Panel (c), Columns 1–4, Table 9). This pattern (which also holds for the level of speed) is also reflected in the relative effect sizes, which is sometimes larger for the spillover than for the direct effect.

suggests that ticketed cars would otherwise have ‘pushed’ the car ahead to drive faster.⁵³ The implications of both of these spillovers are clear-cut: speeding tickets contribute to a drop in speeding in a broader population beyond the ticketed cars. It is not necessarily clear, however, whether these spillovers are all positive from a social welfare perspective (e.g. associated with lower noise and CO2 emissions): ‘backward spillovers’ are responses to cars in front slowing down, which could – in case of sudden braking – increase the risk of accidents.

7 Concluding Summary

Based on unique data that cover the driving histories of 1.2 million cars over several years, we have identified responses to experiencing law enforcement. The results from a regression discontinuity design, which exploit speed level cutoffs with extensive or intensive margin variation in punishment, document that speeding tickets induce a pronounced shift in the speed distribution. The effects are reflected in a decline in the speeding rate by a third and a 70% drop in re-offending. Doubling the fines has only limited additional effects. Event study estimates, which confirm all LATEs from the RDD, further show that these responses are immediate and persistent over several years. Adjustments in speeding are observable in different speed camera zones and there is no evidence of compensatory speeding on un-monitored parts of the road. Instead, the data indicate spillover effects on untreated cars, which also reduce their speed. Given that the WHO (2018) considers effective speed management policies a central strategy to reduce the approximately 1.35 million annual deaths in road traffic crashes, these findings seem relevant.

We present a simple reinforcement learning model which offers a coherent interpretation of the evidence. After being punished, imperfectly informed agents update their priors about the enforcement regime. They learn from law enforcement and adjust their behavior accordingly. Our set-up, which excludes other channels through which past punishment could shape future compliance, is ideal for isolating such learning-induced deterrence effects. The data are consistent with a coarse, discontinuous updating of priors (and reject fine-grained updating). The persistency of the effect further rejects any interpretation of the findings in terms of temporary salience responses of agents with limited attention.

Our results point to the importance of learning and information transmission for mediating deterrence effects. Policy design that aims at facilitating learning effects might therefore constitute an important, and so far under-researched dimension of optimal law enforcement. At the same time, our results allude to the potential benefits of ambiguity in law enforcement. Under coarse updating, an ambiguous enforcement cutoff contributes to larger behavioral responses to punishment. This result might be relevant in numerous domains where offenders face ambiguity about the red line at which an authority’s ‘tolerance’ of illegal behavior ends and enforcement starts (e.g. petty theft, minor drug possession, public nuisances, tax evasion). It is up to future research to examine the domains to which this applies.

⁵³Sensitivity checks indicate that such forward spillovers seem to be larger when the treated car (at, e.g. position $\ell = 2$) faces more scope for speeding (e.g. when the time gap on entering the camera zone is 3–8 (rather than < 5) seconds).

References


- Apel, R. (2013). Sanctions, perceptions, and crime: Implications for criminal deterrence. *Journal of Quantitative Criminology* 29(1), 67–101.
- Ashenfelter, O. and M. Greenstone (2004). Using Mandated Speed Limits to Measure the Value of a Statistical Life. *Journal of Political Economy* 112(S1), 226–267.
- Banerjee, A., E. Duflo, D. Keniston, and N. Singh (2019). The Efficient Deployment of Police Resources: Theory and New Evidence from a Randomized Drunk Driving Crackdown in India. NBER Working Paper No. 26224.
- Barbarino, A. and G. Mastrobuoni (2014). The Incapacitation Effect of Incarceration: Evidence from Several Italian Collective Pardons. *American Economic Journal: Economic Policy* 6(1), 1–37.
- Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci (2015). Self-Confirming Equilibrium and Model Uncertainty. *American Economic Review* 105(2), 646–677.
- Bauernschuster, S. and R. Rekers (2019). Speed Limit Enforcement and Road Safety: Evidence from German Blitzmarathons. Mimeo, University of Passau.
- Bayer, P., R. Hjalmarsson, and D. Pozen (2009). Building Criminal Capital behind Bars: Peer Effects in Juvenile Corrections. *Quarterly Journal of Economics* 124(1), 105–147.
- Becker, G. S. (1968). Crime and Punishment: An Economic Approach. *Journal of Political Economy* 76(2), 169–217.
- Bhuller, M., G. B. Dahl, K. V. Løken, and M. Mogstad (2019). Incarceration, Recidivism and Employment. *Journal of Political Economy*. forthcoming.
- Bérgolo, M., R. Ceni, G. Cruces, M. Giacobasso, and R. Perez-Truglia (2018). Misperceptions about Tax Audits. *AEA Papers and Proceedings* 108, 83–87.
- Calonico, S., M. Cattaneo, M. Farrell, and R. Titiunik (2017). rdrobust: Software for Regression-Discontinuity Designs. *Stata Journal* 17(2), 372–404.
- Calonico, S., M. Cattaneo, and R. Titiunik (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica* 82, 2295–2326.
- Chalfin, A. and J. McCrary (2017). Criminal Deterrence: A Review of the Literature. *Journal of Economic Literature* 55(1), 5–48.
- Chen, M. K. and J. M. Shapiro (2007). Do harsher prison conditions reduce recidivism? A discontinuity-based approach. *American Law and Economics Review* 9(1), 1–29.
- DeAngelo, G. and B. Hansen (2014). Life and Death in the Fast Lane: Police Enforcement and Traffic Fatalities. *American Economic Journal: Economic Policy* 6(2), 231–257.
- DeBacker, J., B. T. Heim, A. Tran, and A. Yuskavage (2015). Legal Enforcement and Corporate Behavior: An Analysis of Tax Aggressiveness after an Audit. *Journal of Law and Economics* 58(2), 291–324.
- Di Tella, R. and E. Schargrodsky (2013). Criminal Recidivism after Prison and Electronic Monitoring. *Journal of Political Economy* 121(1), 28–73.
- Draca, M., S. Machin, and R. Witt (2011). Panic on the Streets of London: Police, Crime, and the July 2005 Terror Attacks. *American Economic Review* 101(5), 2157–2181.
- Drago, F., R. Galbiati, and P. Vertova (2009). The Deterrent Effects of Prison: Evidence from a Natural Experiment. *Journal of Political Economy* 117(2), 257–280.

- Drago, F., R. Galbiati, and P. Vertova (2011). Prison conditions and recidivism. *American Law and Economics Review* 13(1), 103–130.
- Drago, F., F. Mengel, and C. Traxler (2019). Compliance Behavior in Networks: Evidence from a Field Experiment. *American Economic Journal: Applied Economics*. forthcoming.
- Dusek, L., N. Pardo, and C. Traxler (2019). Salience, Incentives and Timely Compliance: Evidence from Speeding Tickets. Working Paper, Hertie School.
- Fudenberg, D. and D. Levine (1993). Self-confirming equilibrium. *Econometrica* 61(3), 523–45.
- Gabaix, X. (2019). Behavioral inattention. In D. Bernheim, S. DellaVigna, and D. Laibson (Eds.), *Handbook of Behavioral Economics*, Volume 2, Chapter 4, pp. 261–343. North-Holland.
- Ganong, P. N. (2012). Criminal rehabilitation, incapacitation, and aging. *American Law and Economics Review* 14(2), 391–424.
- Gehrsitz, M. (2017). Speeding, Punishment, and Recidivism: Evidence from a Regression Discontinuity Design. *Journal of Law and Economics* 60(3), 497–528.
- Glueck, S. (1928). Principles of a Rational Penal Code. *Harvard Law Review* 41(4), 453–482.
- Goncalves, F. and S. Mello (2017). Does the Punishment Fit the Crime? Speeding Fines and Recidivism. SSRN Paper 3064406.
- Goncalves, F. and S. Mello (2018). A Few Bad Apples? Racial Bias in Policing. Working Paper, UCLA.
- Hansen, B. (2015). Punishment and Deterrence: Evidence from Drunk Driving. *American Economic Review* 105(4), 1581–1617.
- Hjalmarsson, R. (2008). Crime and Expected Punishment: Changes in Perceptions at the Age of Criminal Majority. *American Law and Economics Review* 11(1), 209–248.
- Hjalmarsson, R. (2009). Juvenile jails: A path to the straight and narrow or to hardened criminality? *Journal of Law and Economics* 52(4), 779–809.
- Kleven, H. J., M. B. Knudsen, C. T. Kreiner, S. Pedersen, and E. Saez (2011). Unwilling or Unable to Cheat? Evidence From a Tax Audit Experiment in Denmark. *Econometrica* 79(3), 651–692.
- Kuziemko, I. (2013). How should inmates be released from prison? An assessment of parole versus fixed-sentence regimes. *Quarterly Journal of Economics* 128(1), 371–424.
- Lang, M. (2017). Legal Uncertainty as a Welfare Enhancing Screen. *European Economic Review* 91(C), 274–289.
- Lochner, L. (2001). A Theoretical and Empirical Study of Individual Perceptions of the Criminal Justice System. University of Rochester, Working Paper 483.
- Lochner, L. (2007). Individual Perceptions of the Criminal Justice System. *American Economic Review* 97(1), 444–460.
- Makowsky, M. D. and T. Stratmann (2009). Political Economy at Any Speed: What Determines Traffic Citations? *American Economic Review* 99(1), 509–527.
- Mastrobuoni, G. and D. Terlizese (2019). Leave the Door Open? Prison Conditions and Recidivism. Collegio Carlo Albert, Working Paper.
- McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics* 142(2), 698–714.

- Mueller-Smith, M. (2015). The criminal and labor market impacts of incarceration. University of Michigan, Working Paper.
- Nagin, D. S. (2013). Deterrence: A Review of the Evidence by a Criminologist for Economists. *Annual Review of Economics* 5(1), 83–105.
- Nagin, D. S., F. T. Cullen, and C. L. Jonson (2009). Imprisonment and recidivism. *Crime and Justice: A Review of Research* 38(1), 115–200.
- Philippe, A. (2019). Specific Deterrence and Learning. University of Bristol, Working Paper.
- Rincke, J. and C. Traxler (2011). Enforcement Spillovers. *Review of Economics and Statistics* 93(4), 1224–1234.
- Sah, R. K. (1991). Social Osmosis and Patterns of Crime. *Journal of Political Economy* 99(6), 1272–1295.
- Studdert, D. M., S. J. Walter, and J. J. Goldhaber-Fiebert (2017). Once Ticketed, Twice Shy? Specific Deterrence from Road Traffic Laws. Working Paper, Stanford University.
- Traxler, C., F. Westermaier, and A. Wohlschlegel (2018). Bunching on the Autobahn? Speeding responses to a ‘notched’ penalty scheme. *Journal of Public Economics* 157(C), 78–94.
- van Benthem, A. (2015). What is the optimal speed limit on freeways? *Journal of Public Economics* 124(C), 44–62.
- von Liszt, F. (1882). Der Zweckgedanke im Strafrecht (The Idea of Purpose in Criminal Law). In F. von Liszt (Ed.), *Strafrechtliche Aufsätze und Vorträge. Erster Band, 1875–1891*. Berlin: J. Guttentag.
- WHO (2018). *Global Status Report on Road Safety*. Geneva: World Health Organization.
- Yitzhaki, S. (1974). Income tax evasion: A theoretical analysis. *Journal of Public Economics* 3(2), 201–202.


Figures

Figure 1: Updating in response to a speeding ticket




Notes: The figures illustrate possible updating responses to receiving a speeding ticket in period t after a ride at speed s_{t-1} . The figures on the left illustrate the case of fine-grained updating. The figures on the right hand side consider a coarse updating response. The two panels at the top display the adjustment in $q^t(\cdot)$, the corresponding panels at the bottom capture the implications for optimal speed choices.

Figure 2: Share of ‘treated’ rides, enforcement cutoff



Notes: The figure presents the cars' share of *ticketed rides* T_i^1 , i.e. rides after receiving a speeding ticket (relative to all rides in the outcome period), around the enforcement cutoff (1st cutoff). The assignment speed, S_i , is normalized relative to the cutoff (14km/h above the limit). Local linear estimates (with MSE-optimal bandwidth), 95% confidence intervals and mean treatment shares in 0.5km/h-bins, based on car-level observations for the first relevant outcome period (see Section 4.2).


Figure 3: Share of ‘high-fine treated’ rides, high-fine cutoff




Notes: The figure presents the cars' share of *high-fine treated rides* T_i^2 , i.e. rides after receiving a high-fine speeding ticket, around the high-fine cutoff (2nd cutoff). The assignment speed, S_i , is normalized relative to the cutoff (23km/h above the limit). Local linear estimates (with MSE-optimal bandwidth), 95% confidence intervals and mean treatment shares in 0.5km/h-bins, based on car-level observations for the first relevant outcome period (see Section 4.2).

Figure 4: Discontinuities in outcomes at the enforcement cutoff

(a) Outcome: Speeding




(b) Outcome: Speed



Notes: The figures present speeding rates (Panel a) and the cars' mean speed (b), i.e. the cars' average share of rides above the speed limit, around the enforcement cutoff. The assignment speed, S_i , is normalized relative to the cutoff (14km/h above the limit). Local linear estimates (with MSE-optimal bandwidth), 95% confidence intervals and mean outcomes in 0.5km/h-bins, based on car-level observations for the first relevant outcome period.


Figure 5: Change in the speed distribution (enforcement cutoff)




Notes: The figure illustrates speed distributions during the outcome period (with the measured speed normalized relative to the speed limit). The solid green line captures the distribution for all rides from cars with an assignment speed S_i within a 0.5km/h range *below* the enforcement cutoff (i.e., with $D_i^1 = 0$); The dashed, red line plots the distribution for all rides from cars with an assignment speed S_i within a 0.5km/h range *above* the cutoff ($D_i^1 = 1$). The figure does not account for the fuzzy nature of the RDD and thus provides a reduced form (lower bound) indication of the shift in the speed distribution of treated cars.

Figure 6: Discontinuities in outcomes at the high-fine cutoff

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figures present speeding rates (Panel a) and the cars' mean speed (b), i.e. the cars' average share of rides above the speed limit, around the high-fine cutoff. The assignment speed, S_i , is normalized relative to the cutoff (23km/h above the limit). Local linear estimates (with MSE-optimal bandwidth), 95% confidence intervals and mean outcomes in 0.5km/h-bins, based on car-level observations for the first relevant outcome period.

Figure 7: Event study: plot of raw data

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure plots speeding rates (Panel a) and the observed mean speed (b) in weekly intervals before and after receiving the first ticket. The sample includes cars around the first ticket event with a low fine. Cars included in the sample have at least one observation during the pre-ticket period (other than the trigger observation) and at least one observation during the post-ticket period. The blue line (indicated with circles) is based on the raw data. The dark-red line(triangles) excludes the trigger observation from the data.

Figure 8: Event study estimates: responses to a low-fine ticket

(a) Outcome: Speeding




(b) Outcome: Speed



Notes: The figure plots the estimated β_w -coefficients from equ. (9) and their 95% confidence intervals. The dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). (The corresponding estimates are also reported in Columns (1) and (2) of Table A.11.) The sample includes cars around the first ticket event punished by a low fine, for which at least one non-trigger observation before the ticket and at least one observation after the ticket are available. The trigger observation is excluded. Week zero (the last week before receiving the ticket) is the omitted category. Cars: 16,407. Observations: 626,430. Mean speed in week zero: 44.86km/h. Mean speeding rate in week zero: 27%.


Figure 9: Event study: shift in the speed distribution




Notes: The figure depicts speed distributions in the event study sample (the sample used to generate estimates in Figure 8 and Columns (1) and (2) of Table A.11). The speed is normalized relative to the speed limit. The solid green line plots the distribution for all rides made 12 and fewer weeks prior to receiving the 1st speeding ticket, with the trigger observation excluded. The dashed red line plots the distribution for all rides made up to 20 weeks after receiving the ticket. The vertical lines mark the speed limit and the first enforcement cutoff.

Figure 10: Event study estimates: long-run effects

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure plots the estimated β -coefficients and their 95% confidence intervals from an equation analogous to equ. (9), where the single dummies indicate individual months (rather than weeks) before and after receiving a speeding ticket. The dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). (The corresponding estimates are also reported in Columns (1) and (2) of Table A.12.) The sample includes cars around the first ticket event punished by a low fine, for which there exists at least one (non-trigger) observation in each of the three-month intervals before and after the ticket. The trigger observation is excluded. Month zero (the last month before receiving the ticket) is the omitted category. Cars: 4,291. Observations: 991,333. Mean speed in month zero: 44.35km/h. Mean speeding rate in month zero: 26%.

Figure 11: Event study estimates: responses to high- vs low-fine tickets

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure plots the estimated β_w -coefficients from equ. (9) and their 95% confidence intervals for cars receiving either a low- or a high-fine ticket. (The estimates for low-fine tickets replicate Figure 8 from above.) The dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). (The corresponding estimates are also reported in Columns (3) and (4) of Table A.11.) The sample includes cars around the first ticket event punished by a high fine, for which at least one non-trigger observation before the ticket and at least one observation after the ticket are available. The trigger observation is excluded. Week zero is the omitted category. The high-fine sample includes 2,107 cars with 65,606 rides. Mean speed in week zero: 45.75km/h. Mean speeding rate in week zero: 27.9%.

Figure 12: Event study estimates: private owner vs corporation

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure plots the estimated β_w -coefficients from equ. (9) and their 95% confidence intervals, separately for private and corporation cars. The dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). (The corresponding estimates are also reported in Table A.13.) We focus on low-fine tickets and maintain all other sample definitions from above. Week zero (the last week before receiving the ticket) is the omitted category.

Figure 13: Event study estimates: paid vs unpaid tickets

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure plots the estimated β_w -coefficients from equ. (9) and their 95% confidence intervals, separately for cars that paid the fine within 90 days of receiving it and cars that did not. The dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). (The corresponding estimates are also reported in Table A.13.) Week zero (the last week before receiving the ticket) is the omitted category.

Figure 14: Event study estimates by ride sequence

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure plots the coefficients (and their 95% confidence intervals) from a regression analogous to equ. (9) except that week dummies are replaced with indicators for a car's ride sequence. Each dummy indicates five rides in their order before and after receiving the first ticket. Camera-specific calendar month fixed effects are also included. The sample of cars is identical to that used for the main estimates presented in Figure 8. The dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). (The corresponding estimates are also reported in Columns (3) and (4) of Table A.12.) The dummy for the last five rides before the ticket is the omitted category. Cars: 16,414. Observations: 1,171,931. Mean speed/speeding rate during the five rides before the ticket: -3.67km/h (below speed limit) / 29%.

Figure 15: Event study estimates by reoffence pattern

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure plots the estimated β_w -coefficients from equ. (9) and their 95% confidence intervals for two groups: (i) cars that commit (at least) a second offense (eligible for a ticket) after the first ticket and (ii) cars that do not reoffend within our sample period. The dependent variables are the mean speed (Panel a) and the binary speeding indicator (Panel b). The sample includes cars around the first ticket event punished by a low fine, for which at least one non-trigger observation before the ticket and at least one observation after the first ticket are available. (Cases where the second ticket is triggered before the first one is delivered are excluded.) The trigger observation is excluded. Week zero (the last week before receiving the ticket) is the omitted category. (The corresponding estimates are also reported in Columns (1) through (4) of Table A.15.) The sample for the two groups covers, (i) 2,551 reoffending cars with 143,292 rides, observed with a mean speed/speeding rate in week zero of 44.72km/h / 28.3% and (ii) 12,802 non-reoffending cars, with 417,829 rides with a week zero mean speed/speeding rate of 44.12km/h / 24.5%.

Figure 16: Event study estimates for responses to the second ticket

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure plots the estimated β_w -coefficients from to equ. (9) and their 95% confidence intervals. The dependent variables are the mean speed (Panel a) and the binary speeding indicator (Panel b). The sample includes cars around the second ticket event, conditional on previously experiencing a first, low-fine ticket. The sample is restricted to cars for which at least one non-trigger observation before and at least one observation after the second ticket. The trigger observation is excluded. Week zero (the last week before receiving the second ticket) is the omitted category. (The corresponding estimates are also reported in Columns (5) and (6) of Table A.15.) Cars: 1,694. Observations: 101,530. Mean speed/mean speeding rate in week zero: 45.12km/h / 27.7%.

Figure 17: Event study estimates by same/other speed camera zone

(a) Outcome: Speeding



(b) Outcome: Speed



Notes: The figure plots the estimated β_w -coefficients from equ. (9) and their 95% confidence intervals, separately for observations occurring in the same camera zone where the ticket was triggered and in other speed camera zones. The dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). The sample includes cars around the first ticket event punished by a low fine, for which at least one non-trigger observation before the ticket and at least one observation after the ticket are available. The trigger observation is excluded. Week zero (the last week before receiving the ticket) is the omitted category. The estimates are reported in Columns (5) through (8) of Table A.11.

Tables

Table 1: Summary statistics

	‘Not-ticketed’ cars	‘Ticketed’ cars	Total (all cars)
<i>Car characteristics</i>			
Observations (rides)	22,049,809	4,084,958	26,134,767
Number of cars	1,304,791	48,422	1,353,213
Number of tickets	0	56,056	56,056
Observations per car	16.90 (74.84)	84.36 (192.02)	19.31 (82.93)
Driving frequency	2.33 (2.76)	3.06 (2.87)	2.45 (2.79)
Number plate: Local region	0.453 (0.498)	0.455 (0.498)	0.453 (0.498)
Number plate: Prague	0.393 (0.488)	0.439 (0.496)	0.400 (0.490)
<i>Ride characteristics</i>			
Speed	-6.00 (7.73)	-5.17 (8.60)	-5.87 (7.88)
Speeding	0.125 (0.331)	0.189 (0.391)	0.135 (0.342)
(Re)offending	0.000 –	0.015 (0.120)	0.003 (0.051)
Temperature	15.36 (12.19)	14.81 (12.09)	15.27 (12.17)
Windspeed	1.75 (1.46)	1.74 (1.47)	1.75 (1.46)
Hour	12.51 (4.59)	12.65 (4.73)	12.53 (4.61)
Weekend	0.204 (0.403)	0.205 (0.404)	0.204 (0.403)
<i>Ticket/trigger characteristics</i>			
Fine amount (CZK)		1,039 (377)	
Probability of paying the fine		0.933 (0.250)	

Notes: The table reports the number of rides, cars and tickets together with sample means (with standard deviations in parentheses) for cars that did (‘ticketed’) or did not receive a speeding ticket during the sample period (August 2014–2018). Speed indicates the measured speed, relative to the speed limit (in km/h). Number plate distinguishes cars registered in the local region (Central Bohemia, where the municipality of Ricany is located) and Prague. The residual category pools all other regions.

Table 2: Predictions for different ways of updating

	No updating	Fine-grained updating	Coarse updating
Behavioral response to speeding ticket	no response	(small) drop in speed, continued speeding	(large) drop in speed, drop in speeding
Bunching/1 st cutoff (enforcement)	yes (correct prior) no (incorrect prior)	yes (evolving over time)	no
Bunching/2 nd cutoff (low/high fine)	yes ^(a) (correct prior) no (incorrect prior)	yes ^(a) (evolving over time)	no ^(b)
Behavioral response to high- vs low-fine speeding tickets	no (no responses to either)	larger scope for differential effect (reinforcement learning)	limited scope for differential effect

Notes: (a) These predictions implicitly assume that a significant share of drivers have a sufficiently strong taste for (and conditions favouring) speeding, such that they are willing to accept a low-fine speeding ticket. (b) For the case of coarse updating, we assume that the second cutoff is not known. If it were known, we should see bunching (at the second cutoff) under coarse updating, too.

Table 3: Wald estimates for average car: **enforcement cutoff**

	(1) Speeding	(2) (Re)Offending	(3) Speed	(4) Speed ^{p50}	(5) Speed ^{p75}	(6) Speed ^{p90}
Estimate ($\beta^{k=1}$)	-0.0951*** [0.0136]	-0.0051*** [0.0019]	-1.4602*** [0.2774]	-1.3097*** [0.2794]	-1.4972*** [0.2663]	-1.7723*** [0.3032]
Y(left)	0.299	0.007	46.153	46.608	49.678	51.703
Relative effect	-31.80%	-70.31%	-3.16%	-2.81%	-3.01%	-3.43%
Bandwidth	4.483	5.776	4.199	3.871	4.583	4.542

Notes: The table presents Wald estimates for car-level observations at the enforcement cutoff (1st cutoff), more specifically, bias-corrected estimates with MSE-optimal bandwidth and robust standard errors in brackets (Calonico et al., 2014, 2017). The table further indicates the effect size relative to the mean outcome in the 0.5km/h bin below the cutoff, Y(left). Number of observations: 224,816 cars.

Table 4: Wald estimates for average ride: **enforcement and high-fine cutoff**

	(1) Speeding	(2) (Re)Offending	(3) Speed	(4) Speeding	(5) (Re)Offending	(6) Speed
	<i>1st cutoff</i>			<i>2nd cutoff</i>		
Estimate (β^k)	-0.0707*** [0.0139]	-0.0031*** [0.0009]	-0.8804*** [0.3191]	-0.0279 [0.0271]	-0.0025 [0.0034]	-0.8247 [0.6856]
Y(left)	0.253	0.005	44.515	0.216	0.008	44.424
Relative effect	-27.96%	-60.99%	-1.98%	-12.89%	-29.98%	-1.86%
Bandwidth	3.368	3.633	3.718	3.346	2.086	2.844
Obs.	2,505,113	2,505,113	2,505,113	264,587	264,587	264,587

Notes: The table presents Wald estimates for ride-level observations for both the enforcement cutoff (1st cutoff) and the high-fine cutoff (2nd cutoff), more specifically, bias-corrected estimates with MSE-optimal bandwidth and cluster robust standard errors in brackets (Calonico et al., 2014, 2017). Number of observations indicates single rides. Standard errors are clustered at the level of cars (with 224,816 cars in the sample for the first and 16,148 cars for the second cutoff). The table further indicates the effect size relative to the mean outcome in the 0.5km/h bin below the cutoff, Y(left).

Table 5: Wald estimates for subgroups: **enforcement cutoff**

	(1) Infrequent	(2) Frequent	(3) Local region	(4) Prague	(5) Other regions
<i>(A) Outcome: Speeding</i>					
Estimate ($\beta^{k=1}$)	-0.1186*** [0.0244]	-0.0761*** [0.0181]	-0.0953*** [0.0226]	-0.0781*** [0.0262]	-0.0924** [0.0370]
Y(left)	0.328	0.271	0.288	0.295	0.321
Relative effect	-36.21%	-28.04%	-33.06%	-26.48%	-28.81%
Bandwidth	3.535	4.021	4.529	2.275	4.073
<i>(B) Outcome: Mean Speed</i>					
Estimate ($\beta^{k=1}$)	-1.6591*** [0.4994]	-1.2295*** [0.3192]	-1.0293** [0.4150]	-1.8673*** [0.5316]	-1.6526*** [0.6230]
Y(left)	46.619	45.697	45.524	46.238	46.778
Relative effect	-3.56%	-2.69%	-2.26%	-4.04%	-3.53%
Bandwidth	3.510	4.679	5.250	2.230	5.130
Obs.	114,899	109,917	74,638	100,946	49,232

Notes: The table presents subgroup-specific Wald estimates for car-level observations at the enforcement cutoff (1st cutoff), more specifically, bias-corrected estimates with MSE-optimal bandwidth and robust standard errors in brackets (Calonico et al., 2014, 2017). The top panel (A) considers speeding (binary), the lower panel (B) the mean speed outcome (in km/h). Columns (1) and (2) compare infrequent and frequent drivers (according to their average frequency of rides per day, measured during the pre-treatment assignment period), columns (3), (4) and (5) compare cars with number plates from the *Ricany-Region*, from *Prague*, and from *other* regions, respectively. The table further includes the effect size relative to the mean outcome in the 0.5km/h bin below the cutoff, Y(left).

Table 6: Wald estimates for average car: **high-fine cutoff**

	(1)	(2)	(3)	(4)	(5)	(6)
	Speeding	(Re)Offending	Speed	Speed ^{p50}	Speed ^{p75}	Speed ^{p90}
Estimate ($\beta^{k=2}$)	-0.0243 [0.0288]	-0.0058 [0.0104]	-0.7225 [0.7913]	-0.6508 [0.7782]	-0.8824 [0.7895]	-0.6883 [0.7819]
Y(left)	0.258	0.015	45.416	45.789	48.706	50.746
Relative effect	-9.42%	-39.43%	-1.59%	-1.42%	-1.81%	-1.36%
Bandwidth	3.784	2.794	2.793	2.825	3.041	4.013

Notes: The table presents Wald estimates for car-level observations at the high-fine cutoff (2nd cutoff), more specifically, bias-corrected estimates with MSE-optimal bandwidth and robust standard errors in brackets (Calonico et al., 2014, 2017). The table further indicates the effect size relative to the mean outcome in the 0.5km/h bin below the cutoff, $Y(\text{left})$. Number of observations: 16,148 cars.

Table 7: Wald estimates for ‘good’ vs ‘bad’ driving conditions: **high-fine cutoff**

	(1)	(2)	(3)	(4)	(5)	(6)
	Speeding (binary)	Speed (mean)	Speed ^{p90}	Speeding (binary)	Speed (mean)	Speed ^{p90}
Panel A.	<i>Good Conditions</i>			<i>Bad Conditions</i>		
Estimate ($\beta^{k=2}$)	-0.0808* [0.0471]	-1.4711* [0.8681]	-2.0812** [1.0525]	-0.0075 [0.0330]	-0.0809 [0.7750]	-0.5930 [0.8070]
Y(left)	0.381	47.665	53.142	0.176	43.997	48.086
Relative effect	-21.18%	-3.09%	-3.92%	-4.28%	-0.18%	-1.23%
Bandwidth	2.628	2.865	2.409	3.124	2.952	3.273
Obs.	13,446	13,446	13,446	13,639	13,639	13,639
Panel B.	<i>Good Conditions</i>			<i>Bad Conditions</i>		
Estimate ($\beta^{k=2}$)	-0.0796** [0.0393]	-0.8873 [0.6962]	-1.0067 [0.7688]	-0.0038 [0.0270]	-0.3259 [0.5453]	-0.8595 [0.5608]
Y(left)	0.388	47.729	53.754	0.183	44.398	49.492
Relative effect	-20.50%	-1.86%	-1.87%	-2.06%	-0.73%	-1.74%
Bandwidth	3.013	2.940	3.258	3.675	3.743	3.916
Obs.	10,937	10,937	10,937	10,937	10,937	10,937

Notes: The table presents Wald estimates for car-level observations at the high-fine cutoff (2nd cutoff), more specifically, bias-corrected estimates with MSE-optimal bandwidth and robust standard errors in brackets (Calonico et al., 2014, 2017). The table compares the effects on the speeding rate, the mean speed and the p90-speed for rides in good (Columns 1 – 3) and bad driving conditions (Columns 4 – 6). These driving conditions are defined by a median split in the traffic situation: a ride in the outcome period with a minimum time gap of at least 5.84 seconds (the median) to the next car ahead is classified as ‘good condition’ ride. Rides with a time gap of less than 5.84 seconds are considered ‘bad condition’ rides. Panel A presents the estimates for cars observed in either good or bad conditions, i.e. we partially compare different cars. Panel B replicates the estimates for a fixed set of 10,937 cars that are observed under both good and bad traffic conditions. The table also indicates the effect size relative to the mean outcome in the 0.5km/h bin below the cutoff, $Y(\text{left})$.

Table 8: Wald estimates for ‘same’ vs ‘other’ camera zones: **enforcement cutoff**

	(1) same	(2) other	(3) same	(4) other
(A) <i>Outcome: Speeding</i>				
Estimate ($\beta^{k=1}$)	-0.1166*** [0.0184]	-0.0612*** [0.0172]	-0.1065*** [0.0207]	-0.0591*** [0.0170]
Y(left)	0.403	0.192	0.392	0.199
Relative effect	-28.90%	-31.87%	-27.15%	-29.75%
Bandwidth	4.369	3.169	4.029	3.818
(B) <i>Outcome: Mean Speed</i>				
Estimate ($\beta^{k=1}$)	-1.8922*** [0.3716]	-1.0746*** [0.2933]	-1.8827*** [0.3469]	-0.9951*** [0.3452]
Y(left)	47.642	44.225	47.323	44.105
Relative effect	-3.97%	-2.43%	-3.98%	-2.26%
Bandwidth	3.849	5.233	5.040	4.570
Obs.	176,937	166,773	118,894	118,894

Notes: The table presents Wald estimates for car-level observations at the enforcement cutoff (1st cutoff), more specifically, bias-corrected estimates with MSE-optimal bandwidth and robust standard errors in brackets (Calonico et al., 2014, 2017). The top panel (A) considers speeding (binary), the lower panel (B) the mean speed outcome (in km/h). Columns (1) and (3) are based on outcomes measured in the *same* speed camera zone that triggered the assignment speed, S_i . Columns (2) and (4) explore outcomes from *other* zones, i.e. a camera zone that differs from the one where the assignment speed, S_i , was recorded. In columns (3) and (4) the sample is constrained to cars that pass by at the ‘same’ and at least one ‘other’ camera zone; columns (1) and (2) do not condition the sample (i.e. partially compare different cars). The table further indicates the effect size relative to the mean outcome in the 0.5km/h bin below the cutoff, $Y(\text{left})$.

Table 9: Treatment effects and spillovers within lines of cars — outcome: speeding

$\ell =$	First Car Treated					Second Car Treated					Third Car Treated				
$j =$	Car 1	Car 2	Car 3	Car 4	Car 5	Car 1	Car 2	Car 3	Car 4	Car 5	Car 1	Car 2	Car 3	Car 4	Car 5
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
(a) Lines with 2 cars															
Estimate	-0.0576***	-0.0415***				-0.013	-0.0338***								
(β_j^ℓ)	[0.0146]	[0.0114]				[0.0096]	[0.0080]								
Y(left)	0.308	0.203				0.201	0.187								
Rel. effect	-18.725%	-20.460%				-6.461%	-18.089%								
Obs.	458,672					464,764									
(b) Lines with 3 cars															
Estimate	-0.0737***	-0.0546***	-0.0513***			-0.0136	-0.0308**	-0.0036			-0.0079	-0.0240*	-0.0355***		
(β_j^ℓ)	[0.0194]	[0.0157]	[0.0131]			[0.0155]	[0.0129]	[0.0119]			[0.0147]	[0.0139]	[0.0122]		
Y(left)	0.31	0.182	0.142			0.192	0.185	0.137			0.168	0.133	0.132		
Rel. effect	-23.783%	-30.046%	-36.043%			-7.070%	-16.641%	-2.615%			-4.690%	-18.051%	-26.869%		
Obs.	184,535					186,055					187,930				
(c) Lines with 4 cars															
Estimate	-0.0304	-0.0582***	-0.0369*	-0.0105		0.0254	-0.0056	0.0138	0.0218		-0.0150	-0.0046	0.0123	-0.0081	
(β_j^ℓ)	[0.0318]	[0.0216]	[0.0207]	[0.0184]		[0.0231]	[0.0216]	[0.0179]	[0.0210]		[0.0226]	[0.0224]	[0.0239]	[0.0201]	
Y(left)	0.284	0.175	0.136	0.103		0.152	0.133	0.089	0.08		0.188	0.135	0.111	0.101	
Rel. effect	-10.705%	-33.214%	-27.123%	-10.254%		16.720%	-4.221%	15.461%	27.349%		-7.974%	-3.396%	11.094%	-8.016%	
Obs.	84,502					84,969					85,928				
(d) Lines with 5 or more cars															
Estimate	-0.0460	-0.0283	0.0019	-0.0201	-0.0392**	-0.0187	-0.0295	-0.0333*	-0.0114	-0.0143	-0.0269	-0.0053	-0.0504***	-0.0440**	-0.0352**
(β_j^ℓ)	[0.0298]	[0.0225]	[0.0198]	[0.0159]	[0.0160]	[0.0218]	[0.0205]	[0.0172]	[0.0167]	[0.0182]	[0.0240]	[0.0194]	[0.0194]	[0.0182]	[0.0174]
Y(left)	0.296	0.201	0.137	0.101	0.101	0.181	0.159	0.135	0.102	0.109	0.154	0.107	0.137	0.095	0.083
Rel. effect	-15.559%	-14.053%	1.409%	-20.002%	-38.915%	-10.343%	-18.551%	-24.622%	-11.174%	-13.085%	-17.491%	-4.984%	-36.706%	-46.432%	-42.441%
Obs.	94,500					94,509					95,689				

Notes: The table presents Wald estimates (at the level of rides) based on equations (10) and (11). Outcome is speeding (binary). Panels (a) – (d) focus on groups of cars with either two, three, four or five and more cars within a line. The first (second / third) five columns consider cases where the $\ell =$ first (second / third) car within a line has potentially qualified for a speeding ticket (i.e. $D_{\ell g} = \{0, 1\}$). Within each block, the different columns present effects on the $j =$ first, second, ... fifth car within a line. Estimates for cases with $j = \ell$ are marked in bold. All estimates are bias-corrected, with MSE-optimal bandwidth and cluster robust standard errors in brackets (Calonico et al., 2014, 2017). Standard errors are clustered at the level of cars. Number of observations indicates single rides (which is constant within each line/ ℓ^{th} -car-treated ‘block’). Effect size is relative to the mean outcome in the 0.5km/h bin below the cutoff, $Y(\text{left})$.

Table 10: Treatment effects and spillovers within lines of cars — outcome: speed

$\ell =$	First Car Treated					Second Car Treated					Third Car Treated				
$j =$	Car 1	Car 2	Car 3	Car 4	Car 5	Car 1	Car 2	Car 3	Car 4	Car 5	Car 1	Car 2	Car 3	Car 4	Car 5
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
(a) Lines with 2 cars:															
Estimate	-0.5459**	-0.0578				-0.3639*	-0.4023*								
(β_j^ℓ)	[0.2216]	[0.2238]				[0.1987]	[0.2153]								
Y(left)	45.524	44.043				44.486	44.129								
Rel. effect	-1.199%	-0.131%				-0.818%	-0.912%								
Obs.	458,672					464,764									
(b) Lines with 3 cars:															
Estimate	-1.0030***	-0.9437***	-0.7253**			-0.0621	-0.2735	0.0522			-0.4076	-0.1237	-0.1622		
(β_j^ℓ)	[0.3137]	[0.3024]	[0.3212]			[0.3069]	[0.2952]	[0.2694]			[0.2580]	[0.2621]	[0.2497]		
Y(left)	45.782	44.27	43.401			44.41	43.947	43.442			44.259	43.458	43.027		
Rel. effect	-2.191%	-2.132%	-1.671%			-0.140%	-0.622%	0.120%			-0.921%	-0.285%	-0.377%		
Obs.	184,535					186,055					187,930				
(c) Lines with 4 cars:															
Estimate	-0.5704	-0.7003	-0.7472	-0.4332		-0.1976	-0.2975	-0.2395	-0.2214		-0.2098	-0.1156	0.3087	0.7248	
(β_j^ℓ)	[0.4951]	[0.4678]	[0.5325]	[0.4617]		[0.4700]	[0.4551]	[0.4380]	[0.4784]		[0.4045]	[0.3363]	[0.4002]	[0.5063]	
Y(left)	45.039	43.969	43.104	42.586		44.614	43.911	43.394	42.689		43.984	43.389	42.666	42.366	
Rel. effect	-1.267%	-1.593%	-1.734%	-1.017%		-0.443%	-0.677%	-0.552%	-0.519%		-0.477%	-0.266%	0.723%	1.711%	
Obs.	84,502					84,969					85,928				
(d) Lines with 5 or more cars:															
Estimate	-0.3679	-0.3098	0.0587	0.2304	-0.4992	-0.8500**	-1.0043**	-0.4091	-0.3706	-0.4023	-0.5203	-0.6752	-0.8273*	-0.7025	-0.2321
(β_j^ℓ)	[0.4470]	[0.3901]	[0.3736]	[0.4014]	[0.4966]	[0.4197]	[0.4035]	[0.3890]	[0.3989]	[0.3709]	[0.5107]	[0.4503]	[0.4709]	[0.4562]	[0.3887]
Y(left)	45.25	44.034	43.401	42.69	42.468	44.265	43.797	42.823	42.374	41.833	43.9	43.313	42.983	42.57	42.162
Rel. effect	-0.813%	-0.704%	0.135%	0.540%	-1.175%	-1.920%	-2.293%	-0.955%	-0.875%	-0.962%	-1.185%	-1.559%	-1.925%	-1.650%	-0.551%
Obs.	94,500					94,509					95,689				

Notes: The table presents Wald estimates (at the level of rides) based on equations (10) and (11). Outcome is measured speed. Panels (a) – (d) focus on groups of cars with either two, three, four or five and more cars within a line. The first (second / third) five columns consider cases where the $\ell =$ first (second / third) car within a line has potentially qualified for a speeding ticket (i.e., $D_{\ell g} = \{0, 1\}$). Within each block, the different columns present effects on the $j =$ first, second, ... fifth car within a line. Estimates for cases with $j = \ell$ are marked in bold. All estimates are bias-corrected, with MSE-optimal bandwidth and cluster robust standard errors in brackets (Calonico et al., 2014, 2017). Standard errors are clustered at the level of cars. Number of observations indicate single rides (which is constant within each line/ ℓ^{th} -car-treated ‘block’). Effect size is relative to the mean outcome in the 0.5km/h bin below the cutoff, $Y(\text{left})$.

A Additional Figures and Tables


Figure A.1: Photograph of a Speed Camera




Notes: The picture shows a camera at the entry of one speed camera zone.

Figure A.2: Density of running variable around the enforcement cutoff

(a) Distribution ('McCrary Plot')




(b) Histogram




Notes: The figure illustrates the distribution of the assignment speed S_i centered around the enforcement cutoff (14km/h above the speed limit). Panel (a) plots the distribution together with the estimates from McCrary's (2008) heaping test. Panel (b) presents a histogram of the assignment speed S_i over 50 bins (0.2km/h per bin).

Figure A.3: Density of running variable (enforcement cutoff): Evolution over time I (all cars)




Notes: The figure plots the distribution of observed speed measures – centered around the enforcement cutoff (14km/h above the speed limit) – in six semi-annual intervals, starting with the first month the radars were operating. Sample includes all cars. Estimates are from McCrary’s (2008) heaping test.

Figure A.4: Density of running variable (enforcement cutoff): Evolution II (regional cars)



Notes: The figure plots the distribution of observe speed measures – centered around the enforcement cutoff (14km/h above the speed limit) – in six semi-annual intervals, starting with the first month the radars were operating. Sample includes only ‘regional cars’, defined as cars with above-median driving frequency and a number plate from the region. Estimates are from McCrary’s (2008) heaping test.


Figure A.5: Continuity of other characteristics around the enforcement cutoff




Notes: The figures document continuity for several observables around the enforcement cutoff. The figures depict the cars' number of rides as well as the driving frequency during the assignment period (pre-treatment); the traffic density (the logged time difference to the car ahead, measured for the ride with the maximum speed); indicators for the cars' number plate (Prague and local Region) and trigger rides on the weekend; indicators for trigger rides in the morning, afternoon, and at night, respectively. Further variables are considered in Table A.1.

Figure A.6: Density of running variable around the high-fine cutoff

(a) Distribution ('McCrary Plot')




(b) Histogram




Notes: The figure illustrates the distribution of the assignment speed S_i centered around the high-fine cutoff (23km/h above the speed limit). Panel (a) plots the distribution together with the estimates from McCrary's (2008) heaping test. Panel (b) presents a histogram of the assignment speed S_i over 50 bins (0.2km/h per bin).

Figure A.7: Share of ‘ticketed’ rides around the high-fine cutoff



Notes: The figure presents the cars’ share of *ticketed rides* T_i^1 , i.e., rides after receiving any speeding ticket (relative to all rides in the outcome period), around the *high-fine cutoff* (2nd cutoff). If high-fine tickets are sent out more quickly than low-fine tickets, this could produce a discontinuity in T_i^1 at the second cutoff. This is not supported by the data. The assignment speed, S_i , is normalized relative to the high-fine cutoff (23km/h above the limit). Local linear estimates (with a MSE-optimal bandwidth), 95% confidence intervals and mean treatment shares in 0.5km/h-bins, based on car-level observations.


Figure A.8: Relative change in speed distribution (enforcement cutoff)




Notes: The figure plots the relative change in the mass between the two speed distributions from Figure 5 – the one for cars with an assignment speed S_i in a 0.5km/h range *above* and the other for cars with an assignment speed in a 0.5km/h range *below* the enforcement – in 7km/h bins above and below the actual speed limit. In computing the relative changes, we first normalized the observed mass in the groups marginally below and above the cutoff. Recall that, among the latter group, only 80% are actually treated. In addition, the pre-treatment speed distributions (i.e., for rides observed during the assignment period) are, by construct, different between the two groups. The relative changes thus represent approximations for the true treatment impact on the speed distribution.

Figure A.9: Sensitivity of reduced form estimates at the enforcement cutoff

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure depicts reduced form estimates (with 95% confidence intervals) for car-level observations at the enforcement cutoff, varying the bandwidth in 0.5km/h steps from 0.5 to 8.0km/h of assignment speed S_i . The outcome is the speeding (top panel) and the mean speed (lower panel). The different coefficients should be compared with the results reported in Table A.3.

Figure A.10: Sensitivity of *car-level* estimates at the enforcement cutoff

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure depicts Wald estimates at the car-level, with 95% CI for the enforcement cutoff (1st cutoff) for different assignment (*a*, in months) and follow-up periods (*f*). Panel (a) presents outcomes for or the speeding rate, panel (b) the mean speed (in km/h). Number of observations, bandwidth, and relative effect sizes are reported in Tables A.6 and A.7, respectively.

Figure A.11: Sensitivity of *ride-level* estimates at the enforcement cutoff

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure depicts Wald estimates at the *ride-level*, with 95% CI for the enforcement cutoff (1st cutoff) for different assignment (a , in months) and follow-up periods (f). Panel (a) presents outcomes for the speeding rate, panel (b) for the mean speed (in km/h).

Figure A.12: Sensitivity of reduced form estimates at the high-fine cutoff

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure depicts reduced form estimates (with 95% confidence intervals) for car-level observations at the high-fine cutoff, varying the bandwidth in 0.5km/h steps from 0.5 to 8.0km/h of assignment speed S_i . Outcome: Speed (in km/h). Panel (a) presents outcomes for the speeding rate, panel (b) for the mean speed (in km/h).

Figure A.13: Sensitivity of *car-level* estimates at the high-fine cutoff

(a) Outcome: Speeding




(b) Outcome: Speed



Notes: The figure depicts Wald estimates (at the car-level) with 95% CI for the high-fine cutoff (2nd cutoff) for different assignment (a , in months) and follow-up periods (f). Outcome: Speed (in km/h). Observations pooled at level of cars.


Figure A.14: Updating in response to low- vs high-fine speeding ticket




Notes: Following the structure of the lower two panels of Fig. 1, the figure illustrate coarse updating and optimal speeding responses to receiving a speeding ticket with either a low ($\frac{\partial q_L^t}{\partial s}$) or a high fine ($\frac{\partial q_H^t}{\partial s}$). For the depicted case, where a higher fine increases the convexity of the marginal expected costs (relative to the updating response followed a low-fine ticket), we would observe pronounced differences in speed choices under ‘good’ driving conditions (like those in c_{t-1}^{good} , which ultimately triggered a high-fine ticket). For driving conditions that are less favourable (such as for c_t^{bad}), however, the differentially updated (marginal) cost functions would imply almost indistinguishable speed choices.

Figure A.15: Responses to high- vs. low-fine ticket by driving conditions

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure plots the estimated β_w -coefficients from equ. (9) and their 95% confidence intervals for cars observed in relatively good or relatively bad driving conditions, receiving either a low- or high-fine ticket. Driving conditions are defined analogously to the RDD analysis (see Table 7): a ride in the outcome period with a minimum time gap of at least [less than] 5.84sec to the next car ahead is classified as 'good condition' ['bad condition'] ride. The dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). The sample includes cars around the first ticket event punished either by a low or a high fine, for which at least one non-trigger observation before the ticket and at least one observation after the ticket are available. Moreover, similar to Panel B of Table 7, we focus on cars that are observed in both good and bad traffic conditions. The trigger observation is excluded. Week zero is the omitted category. The high-fine sample includes 1,275 cars with 25,850 rides in good and 30,733 rides in bad conditions. The low-fine sample has 11,020 cars with 269,091 rides in good and 284,147 in bad driving conditions.

Figure A.16: Heterogeneity: speed and speeding responses by driving frequency

(a) Outcome: Speeding




(b) Outcome: Speed




Notes: The figure plots the estimated β_w -coefficients from equ. (9) and their 95% confidence intervals. The dependent variables are the speeding dummy (Panel a) and measured speed s_{it} in km/h(b). The sample is divided into infrequent and frequent cars, as defined by a median split according to the average daily pre-treatment rides (i.e., rides during the period between a car's first appearance and the day it receives the first ticket). We focus on low-fine tickets and maintain all other sample definitions from above. Week zero (the last week before receiving the ticket) is the omitted category. The corresponding estimates are also reported in Tables A.14.

Figure A.17: Heterogeneity: speed and speeding responses by number plate region

(a) Outcome: Speeding



(b) Outcome: Speed



Notes: The figure plots the estimated β_w -coefficients from equ. (9) and their 95% confidence intervals. The dependent variables are the speeding dummy (Panel a) and measured speed s_{it} in km/h (b). The sample is split by number plate into cars from the 'Local' region (i.e. where the municipality of Ricany is located), 'Prague', and all 'Other' regions. We focus on low-fine tickets and maintain all other sample definitions from above. Week zero (the last week before receiving the ticket) is the omitted category. The estimates are also reported in Table A.14.

Table A.1: Balancing checks: enforcement cutoff

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	
	driving frequ.	# rides (pre-treat)	no. plate: Prague	no. plate: Local region	traffic	temp	wind	hour (trigger ride)			
								6–9am	9–12pm	12–3pm	
Estimate	−0.0048 [0.0135]	−0.1127 [0.8298]	0.0384* [0.0204]	0.0182 [0.0206]	0.0915 [0.0587]	−0.7626** [0.3722]	−0.1172 [0.0819]	0.0211 [0.0167]	−0.0306 [0.0220]	0.0139 [0.0214]	
Y(left)	0.259	10.645	0.456	0.302	2.926	12.010	1.833	0.141	0.217	0.206	
	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	
	hour (trigger ride)			day of week (trigger ride)							month
	3–6pm	6–9pm	9–12am	Tue	Wed	Thu	Fr	Sa	Su	02	
Estimate	0.0108 [0.0195]	0.0167 [0.0167]	0.0067 [0.0114]	−0.0173 [0.0175]	0.0081 [0.0153]	0.0215 [0.0171]	0.0058 [0.0140]	0.0122 [0.0166]	−0.0029 [0.0187]	0.0192* [0.0115]	
Y(left)	0.187	0.138	0.048	0.123	0.120	0.121	0.128	0.184	0.194	0.047	
	(21)	(22)	(23)	(24)	(25)	(26)	(27)	(28)	(29)	(30)	
	month of year (trigger ride)										
	03	04	05	06	07	08	09	10	11	12	
Estimate	0.0164 [0.0145]	0.0151 [0.0137]	−0.0173 [0.0152]	−0.0319** [0.0132]	0.0098 [0.0144]	0.0018 [0.0114]	0.0192* [0.0101]	−0.0039 [0.0113]	−0.0032 [0.0109]	−0.0087 [0.0151]	
Y(left)	0.098	0.071	0.103	0.124	0.087	0.065	0.042	0.080	0.083	0.100	

Table A.2: Balancing checks: high-fine cutoff

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	
	driving frequ.	# rides (pre-treat)	no. plate: Prague	no. plate: Local region	traffic	temp	wind	hour (trigger ride)			
								6–9am	9–12pm	12–3pm	
Estimate	0.0024 [0.0230]	1.0433 [1.5120]	−0.0151 [0.0422]	0.0792* [0.0450]	0.0499 [0.0993]	−0.1306 [0.9319]	0.0948 [0.1520]	0.0413 [0.0375]	0.0278 [0.0334]	0.0106 [0.0403]	
Y(left)	0.240	8.481	0.481	0.303	3.293	11.125	1.649	0.141	0.174	0.162	
	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	
	hour (trigger ride)			day of week (trigger ride)							month
	3–6pm	6–9pm	9–12am	Tue	Wed	Thu	Fr	Sa	Su	02	
Estimate	−0.0381 [0.0355]	0.0148 [0.0321]	−0.0070 [0.0264]	0.0000 [0.0264]	0.0070 [0.0321]	−0.0246 [0.0333]	0.0274 [0.0260]	−0.0574 [0.0412]	0.0170 [0.0447]	0.0422** [0.0174]	
Y(left)	0.199	0.145	0.095	0.104	0.129	0.116	0.083	0.241	0.241	0.029	
	(21)	(22)	(23)	(24)	(25)	(26)	(27)	(28)	(29)	(30)	
	month of year (trigger ride)										
	03	04	05	06	07	08	09	10	11	12	
Estimate	0.0028 [0.0226]	0.0106 [0.0231]	0.0105 [0.0278]	−0.0659** [0.0307]	−0.0035 [0.0311]	−0.0358 [0.0258]	0.0428** [0.0210]	−0.0174 [0.0242]	−0.0200 [0.0210]	−0.0030 [0.0327]	
Y(left)	0.066	0.058	0.141	0.141	0.104	0.075	0.025	0.087	0.083	0.124	

Notes: Tables A.1 and A.2 present a series of balancing checks for the enforcement (A.1) and the high-fine cutoff (A.2), respectively. The tables report bias-corrected RD estimates with MSE-optimal bandwidth and robust standard errors in brackets (Calonico et al., 2014, 2017). The variables capture the cars' driving frequency and number of rides during the assignment period (pre-treatment), indicators for the cars' number plate (residual category: other regions), driving conditions (temperature, wind, and traffic density, measured by the logged time difference to the next car in front) as well as time and date information for the 'trigger ride'. Y(left) indicates the mean of the dependent variable in the 0.5km/h bin below the respective cutoff. Number of observations for all specifications from Tab. A.1: 224,816 cars. Tab. A.2: 16,148 cars.

Table A.3: Reduced form estimates: enforcement cutoff

	(1) Ticketed	(2) Speeding	(3) (Re)Offending	(4) Speed	(5) Speed ^{p90}
Estimate (δ, τ)	0.7866*** [0.0127]	-0.0812*** [0.0146]	-0.0044** [0.0019]	-1.3512*** [0.2814]	-1.4023*** [0.2711]
Y(left)	0.017	0.299	0.007	46.153	51.703
Bandwidth	2.428	2.228	2.619	2.270	3.353

Notes: The table presents reduced form results for the enforcement cutoff: bias-corrected RD estimates with MSE-optimal bandwidth and robust standard errors in brackets (Calonico et al., 2014, 2017). Number of observations: 224,816 cars.

Table A.4: Reduced form estimates: high-fine cutoff

	(1) High-fine Treated	(2) Ticketed (2 nd cutoff)	(3) Speeding	(4) (Re)Offending	(5) Speed	(6) Speed ^{p90}
Estimate (δ, τ)	0.8145*** [0.0189]	0.0117 [0.0350]	-0.0199 [0.0235]	-0.0048 [0.0085]	-0.5938 [0.6494]	-0.5632 [0.6381]
Y(left)	0.808	0.808	0.258	0.015	45.416	50.746
Bandwidth	5.080	2.649	3.784	2.794	2.793	4.013

Notes: The table presents reduced form results for the enforcement cutoff: bias-corrected RD estimates with MSE-optimal bandwidth and robust standard errors in brackets (Calonico et al., 2014, 2017). Number of observations: 16,148 cars.

Table A.5: Reduced form estimates: driving responses at enforcement and high-fine cutoff

	(1) Rides (count)	(2) Ever-return (binary)	(3) Rides (count)	(4) Ever-return (binary)
	<i>1st cutoff</i>		<i>2nd cutoff</i>	
Estimate (τ)	0.8812 [0.6501]	0.0389** [0.0173]	-0.2831 [1.6193]	0.0022 [0.0382]
Y(left)	7.263	0.509	7.420	0.557
Bandwidth	2.710	2.293	2.589	2.661
Obs. (Cars)	465,518	465,518	27,774	27,774

Notes: The table presents reduced form results examining extensive margin driving responses (e.g. avoiding roads with speed cameras) at the enforcement (Columns 1 – 2) and the high-fine (Columns 3 – 4) cutoffs, respectively. The samples are defined as in our main estimates (see Tables 3 and 6 and the sample definitions discussed in Section 4.1) but also includes cars that were *not* observed during the outcome period. The dependent variables measure the number of rides during the outcome period (Columns 1 and 3) or indicate whether a car observed during the assignment period was ever observed again (Columns 2 and 4). Bias-corrected RD estimates based on car-level observations with MSE-optimal bandwidth and robust standard errors in brackets (Calonico et al., 2014, 2017). Y(left) indicates the mean outcome in the 0.5km/h bin below the cutoff.

Table A.6: Car-level estimates for varying a and f parameters – Outcome: Speed

		f = 3	f = 4	f = 5	f = 6
Estimate (β) Y(left) Relative effect Bandwidth Obs	$a=3$	-1.6654*** [0.4073]	-1.4110*** [0.3583]	-1.3028*** [0.3017]	-1.7686*** [0.3438]
		46.273	46.205	46.238	46.287
		-3.599	-3.054	-2.818	-3.821
		2.580	3.042	3.763	2.439
		203,459	221,303	235,072	246,178
Estimate (β) Y(left) Relative effect Bandwidth Obs	$a=4$	-1.4734*** [0.3222]	-1.4602*** [0.2774]	-1.4402*** [0.2726]	-1.7563*** [0.3133]
		46.225	46.153	46.204	46.243
		-3.187	-3.164	-3.117	-3.798
		3.423	4.199	3.774	2.515
		207,737	224,816	239,220	250,887
Estimate (β) Y(left) Relative effect Bandwidth Obs	$a=5$	-1.3400*** [0.3133]	-1.3793*** [0.2919]	-1.4013*** [0.2768]	-1.2784*** [0.2478]
		46.148	46.083	46.121	46.160
		-2.904	-2.993	-3.038	-2.769
		3.017	3.150	3.074	3.436
		211,031	227,732	241,375	253,579
Estimate (β) Y(left) Relative effect Bandwidth Obs	$a=6$	-1.2608*** [0.3120]	-1.2932*** [0.2837]	-1.3481*** [0.2862]	-1.1707*** [0.2597]
		46.033	46.006	46.000	45.967
		-2.739	-2.811	-2.931	-2.547
		2.915	3.136	2.670	2.913
		213,471	230,170	243,558	255,215

Notes: The table presents Wald estimates for car-level observations at the enforcement cutoff, considering different combinations of assignment (a , in months) and follow-up periods (f). Outcome: Speed (in km/h). The table includes the relative effect size (relative to the mean outcome in the 0.5km/h bin below the cutoff, $Y(\text{left})$), the MSE-optimal bandwidths, and the number of observations (cars) with S_i in the $[-14, 9]$ km/H range around the cutoff. Illustration of point estimates provided in Fig. A.10.

Table A.7: Car-level estimates for varying a and f parameters – Outcome: Speeding

		f =3	f =4	f =5	f =6
Estimate (β)	$a=3$	-0.1145*** [0.0231]	-0.0912*** [0.0188]	-0.0916*** [0.0166]	-0.1066*** [0.0182]
Y(left)		0.314	0.299	0.301	0.305
Relative effect		-36.490	-30.456	-30.437	-34.964
Bandwidth		2.285	2.795	3.107	2.309
Obs		203,459	221,303	235,072	246,178
Estimate (β)	$a=4$	-0.0933*** [0.0176]	-0.0951*** [0.0136]	-0.1017*** [0.0139]	-0.1066*** [0.0163]
Y(left)		0.310	0.299	0.302	0.305
Relative effect		-30.116	-31.798	-33.735	-34.990
Bandwidth		3.251	4.483	3.846	2.560
Obs		207,737	224,816	239,220	250,887
Estimate (β)	$a=5$	-0.1042*** [0.0178]	-0.1060*** [0.0158]	-0.1081*** [0.0148]	-0.1025*** [0.0132]
Y(left)		0.310	0.301	0.302	0.306
Relative effect		-33.589	-35.214	-35.754	-33.537
Bandwidth		2.725	2.879	2.911	3.390
Obs		211,031	227,732	241,375	253,579
Estimate (β)	$a=6$	-0.1090*** [0.0170]	-0.1095*** [0.0150]	-0.1114*** [0.0146]	-0.1055*** [0.0136]
Y(left)		0.312	0.304	0.304	0.305
Relative effect		-34.953	-35.980	-36.593	-34.592
Bandwidth		2.746	2.983	2.771	3.036
Obs		213,471	230,170	243,558	255,215

Notes: The table presents Wald estimates for car-level observations at the enforcement cutoff, considering different combinations of assignment (a , in months) and follow-up periods (f). Outcome: speeding (binary). The table includes the relative effect size (relative to the mean outcome in the 0.5km/h bin below the cutoff, $Y(\text{left})$), the MSE-optimal bandwidths, and the number of observations (cars) with S_i in the $[-14, 9]$ km/H range around the cutoff. Illustration of point estimates provided in Fig. A.10.

Table A.8: Wald estimates for ‘good’ vs ‘bad’ driving conditions: enforcement cutoff

	(1)	(2)	(3)	(4)	(5)	(6)
	Speeding (binary)	Speed (mean)	Speed ^{p90}	Speeding (binary)	Speed (mean)	Speed ^{p90}
Panel A.	<i>Good Conditions</i>			<i>Bad Conditions</i>		
Estimate ($\beta^{k=1}$)	-0.1512*** [0.0218]	-2.4135*** [0.4252]	-2.7780*** [0.4723]	-0.0536*** [0.0133]	-0.7510*** [0.2765]	-0.7596** [0.3388]
Y(left)	0.410	47.770	52.828	0.197	44.390	48.744
Relative effect	-36.86%	-5.05%	-5.26%	-27.25%	-1.69%	-1.56%
Bandwidth	2.962	2.951	2.987	4.904	5.020	3.717
Obs.	171,329	171,329	171,329	185,829	185,829	185,829
Panel B.	<i>Good Conditions</i>			<i>Bad Conditions</i>		
Estimate ($\beta^{k=1}$)	-0.1283*** [0.0149]	-1.8548*** [0.3454]	-2.4969*** [0.3902]	-0.0533*** [0.0130]	-0.8609*** [0.2835]	-0.8842*** [0.3099]
Y(left)	0.395	47.219	53.062	0.190	44.386	49.528
Relative effect	-32.46%	-3.93%	-4.71%	-28.37%	-1.94%	-1.78%
Bandwidth	6.925	5.255	5.207	5.193	4.844	4.394
Obs.	132,342	132,342	132,342	132,342	132,342	132,342

Notes: The table presents Wald estimates for car-level observations at the enforcement cutoff, more specifically, bias-corrected estimates with MSE-optimal bandwidth and robust standard errors in brackets (Calonico et al., 2014, 2017). The table compares the effects on the speeding rate, the mean speed and the p90-speed for riders in good (Columns 1 – 3) and bad driving conditions (Columns 4 – 6). These driving conditions are defined by a median split in the traffic situation of rides. More specifically, a ride in the outcome period with a minimum time gap of at least 5.78 seconds (the median) to the next car ahead is classified as ‘good condition’ ride. Rides with a time gap of less than 5.78 seconds are considered ‘bad condition’ rides. Panel A presents the estimates for cars observed in either good or bad conditions (i.e. we partially compare different cars). Panel B replicates the estimates for a fixed set of 132,342 cars that are observed in both good and bad traffic conditions. The table also indicates the effect size relative to the mean outcome in the 0.5km/h bin below the cutoff, $Y(\text{left})$.

Table A.9: Wald estimates for subgroups: high-fine cutoff

	(1)	(2)	(3)	(4)	(5)
	Infrequent	Frequent	Local region	Prague	Other regions
(A) <i>Outcome: Speeding</i>					
Estimate ($\beta^{k=2}$)	-0.0510 [0.0432]	-0.0113 [0.0425]	-0.0732 [0.0480]	0.0229 [0.0451]	-0.0644 [0.0834]
Y(left)	0.289	0.220	0.279	0.251	0.244
Relative effect	-17.67%	-5.13%	-26.18%	9.13%	-26.38%
Bandwidth	3.766	2.748	3.798	3.108	2.658
(B) <i>Outcome: Mean Speed</i>					
Estimate ($\beta^{k=2}$)	-0.9985 [0.9547]	-0.5009 [1.2346]	0.3482 [1.2142]	-0.5939 [0.9385]	-2.1516 [2.2243]
Y(left)	46.627	43.924	44.839	46.075	44.757
Relative effect	-2.14%	-1.14%	0.78%	-1.29%	-4.81%
Bandwidth	3.223	2.697	4.044	3.177	2.829
Obs.	8,075	8,073	5,310	7,607	3,231

Notes: The table presents subgroup-specific Wald estimates for car-level observations at the high-fine cutoff, more specifically, bias-corrected estimates with MSE-optimal bandwidth and robust standard errors in brackets (Calonico et al., 2014, 2017). The top panel (A) considers speeding (binary), the lower panel (B) the mean speed (in km/h). Columns (1) and (2) compare infrequent and frequent drivers (according to their average frequency of rides per day, measured during the pre-treatment assignment period), columns (3), (4) and (5) compare cars with number plates from the *Ricany-Region*, from *Prague*, and from *other* regions, respectively. The table further includes the effect size relative to the mean outcome in the 0.5km/h bin below the cutoff, $Y(\text{left})$.

Table A.10: Driving outcomes on unmonitored road: enforcement cutoff

	<i>car-level</i>		<i>ride-level</i>	
	(1)	(2)	(3)	(4)
	Speed	Speed ⁹⁰	Speed	log(time)
Estimate (τ)	-1.1174 [1.8909]	-3.9075* [2.0315]	-0.6468 [0.4417]	0.0645** [0.0323]
Y(left)	28.515	37.431	26.230	5.138
Relative effect	-3.92%	-10.44%	-2.47%	1.26%
Bandwidth	3.128	2.744	1.984	1.556
Obs.	3,683	3,683	88,596	88,596

Notes: The table presents reduced form results examining driving responses on the un-monitored road. The table reports estimates at the car level (Columns 1–2) as well as at the ride level (Columns 3–4). In the former, the dependent variables are mean speed and the 90th-percentile speed on the un-monitored road (collapsed at the car level), respectively. The latter specifications use the speed per ride, s_{it} , as well as the log of the travel time between the radars. As in our main RDD estimates, the sample includes all cars that are observed on a trip that passed the two speed cameras during their first outcome period after recording an assignment speed in the range around the enforcement cutoff. Bias-corrected RD estimates with MSE-optimal bandwidth and robust (Columns 1–2) and clustered (3–4) standard errors in brackets (Calonico et al., 2014, 2017). $Y(\text{left})$ indicates the mean outcome in the 0.5km/h bin below the cutoff, which is used to compute the relative effect size.

Table A.11: Event analysis: responses to the first ticket

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	<i>Low Fine Ticket</i>		<i>High Fine Ticket</i>		<i>Same Zone</i>		<i>Other Zone</i>	
	Speeding	Speed	Speeding	Speed	Speeding	Speed	Speeding	Speed
Week -11 before	0.002 (0.004)	0.045 (0.087)	-0.003 (0.014)	-0.415 (0.290)	0.007 (0.007)	-0.093 (0.143)	0.004 (0.005)	0.201* (0.109)
Week -10 before	0.001 (0.004)	-0.07 (0.089)	-0.003 (0.013)	-0.431* (0.254)	0.004 (0.007)	-0.269* (0.151)	0.005 (0.005)	0.155 (0.105)
Week -9 before	0.004 (0.004)	-0.034 (0.086)	-0.002 (0.014)	-0.656** (0.272)	0.009 (0.007)	-0.010 (0.141)	0.006 (0.005)	0.039 (0.106)
Week -8 before	0.001 (0.004)	-0.099 (0.086)	-0.016 (0.013)	-0.411 (0.254)	0.002 (0.007)	-0.256* (0.144)	0.005 (0.005)	0.100 (0.104)
Week -7 before	-0.004 (0.004)	-0.151* (0.086)	0.007 (0.013)	-0.664*** (0.249)	0.003 (0.007)	-0.269* (0.141)	-0.003 (0.005)	0.027 (0.104)
Week -6 before	0.003 (0.004)	-0.073 (0.084)	-0.006 (0.012)	-0.525** (0.243)	0.004 (0.007)	-0.123 (0.136)	0.005 (0.005)	-0.008 (0.105)
Week -5 before	0.002 (0.004)	-0.1 (0.082)	-0.016 (0.013)	-0.466* (0.276)	-0.002 (0.007)	-0.053 (0.131)	0.005 (0.005)	-0.093 (0.101)
Week -4 before	0.010*** (0.004)	-0.03 (0.081)	-0.016 (0.012)	-0.774*** (0.249)	0.010 (0.007)	-0.002 (0.129)	0.011** (0.005)	-0.015 (0.102)
Week -3 before	0.008* (0.004)	0.028 (0.080)	-0.002 (0.012)	-0.343 (0.241)	0.009 (0.007)	0.034 (0.126)	0.008* (0.005)	0.054 (0.100)
Week -2 before	0.007* (0.004)	0.061 (0.078)	-0.008 (0.012)	-0.348 (0.235)	0.009 (0.006)	0.026 (0.126)	0.008* (0.005)	0.124 (0.095)
Week -1 before	0.008** (0.004)	0.133* (0.077)	0.003 (0.011)	-0.465** (0.215)	0.009 (0.006)	0.104 (0.123)	0.008* (0.005)	0.173* (0.095)
Week 1 after	-0.074*** (0.004)	-0.972*** (0.076)	-0.073*** (0.011)	-1.512*** (0.229)	-0.102*** (0.006)	-1.394*** (0.121)	-0.053*** (0.004)	-0.640*** (0.096)
Week 2 after	-0.090*** (0.004)	-1.162*** (0.077)	-0.091*** (0.012)	-1.628*** (0.247)	-0.125*** (0.006)	-1.620*** (0.124)	-0.064*** (0.004)	-0.803*** (0.095)
Week 3 after	-0.097*** (0.004)	-1.230*** (0.078)	-0.095*** (0.012)	-2.006*** (0.244)	-0.139*** (0.007)	-1.714*** (0.124)	-0.066*** (0.004)	-0.840*** (0.096)
Week 4 after	-0.095*** (0.004)	-1.252*** (0.082)	-0.102*** (0.012)	-1.834*** (0.233)	-0.137*** (0.007)	-1.765*** (0.131)	-0.064*** (0.005)	-0.868*** (0.100)
Week 5 after	-0.095*** (0.004)	-1.365*** (0.081)	-0.096*** (0.012)	-1.805*** (0.235)	-0.136*** (0.007)	-2.003*** (0.129)	-0.069*** (0.005)	-0.889*** (0.098)
Week 6 after	-0.097*** (0.004)	-1.263*** (0.083)	-0.093*** (0.013)	-1.828*** (0.257)	-0.136*** (0.007)	-1.780*** (0.133)	-0.070*** (0.005)	-0.841*** (0.099)
Week 7 after	-0.096*** (0.004)	-1.302*** (0.081)	-0.092*** (0.013)	-1.564*** (0.231)	-0.145*** (0.007)	-1.989*** (0.133)	-0.063*** (0.005)	-0.796*** (0.100)
Week 8 after	-0.095*** (0.004)	-1.224*** (0.086)	-0.107*** (0.012)	-1.900*** (0.241)	-0.141*** (0.007)	-1.885*** (0.140)	-0.066*** (0.005)	-0.744*** (0.103)
Week 9 after	-0.097*** (0.004)	-1.242*** (0.082)	-0.103*** (0.012)	-2.000*** (0.249)	-0.141*** (0.007)	-1.835*** (0.133)	-0.068*** (0.005)	-0.812*** (0.100)
Week 10 after	-0.100*** (0.004)	-1.351*** (0.084)	-0.098*** (0.013)	-1.838*** (0.254)	-0.140*** (0.007)	-1.924*** (0.135)	-0.073*** (0.005)	-0.932*** (0.101)
Week 11 after	-0.099*** (0.004)	-1.257*** (0.083)	-0.100*** (0.012)	-1.591*** (0.235)	-0.142*** (0.007)	-1.867*** (0.132)	-0.068*** (0.005)	-0.809*** (0.103)
Week 12 after	-0.099*** (0.004)	-1.477*** (0.086)	-0.109*** (0.013)	-1.767*** (0.248)	-0.150*** (0.007)	-2.028*** (0.139)	-0.062*** (0.005)	-1.045*** (0.106)
Week 13 after	-0.098*** (0.004)	-1.250*** (0.085)	-0.110*** (0.013)	-1.755*** (0.256)	-0.141*** (0.007)	-1.858*** (0.137)	-0.070*** (0.005)	-0.840*** (0.105)
Week 14 after	-0.102*** (0.004)	-1.283*** (0.085)	-0.116*** (0.012)	-1.938*** (0.256)	-0.146*** (0.007)	-1.819*** (0.135)	-0.071*** (0.005)	-0.890*** (0.104)
Week 15 after	-0.099*** (0.004)	-1.349*** (0.085)	-0.111*** (0.013)	-1.956*** (0.253)	-0.143*** (0.007)	-1.999*** (0.135)	-0.068*** (0.005)	-0.855*** (0.102)
Week 16 after	-0.109*** (0.004)	-1.447*** (0.084)	-0.110*** (0.014)	-1.972*** (0.272)	-0.156*** (0.007)	-2.120*** (0.135)	-0.075*** (0.005)	-0.892*** (0.104)
Week 17 after	-0.101*** (0.004)	-1.345*** (0.087)	-0.108*** (0.013)	-1.922*** (0.251)	-0.149*** (0.007)	-2.085*** (0.141)	-0.065*** (0.005)	-0.769*** (0.103)
Week 18 after	-0.108*** (0.004)	-1.370*** (0.086)	-0.093*** (0.013)	-1.983*** (0.265)	-0.155*** (0.007)	-2.032*** (0.134)	-0.075*** (0.005)	-0.880*** (0.107)
Week 19 after	-0.110*** (0.004)	-1.509*** (0.086)	-0.112*** (0.013)	-1.904*** (0.272)	-0.159*** (0.007)	-2.239*** (0.137)	-0.078*** (0.005)	-0.996*** (0.106)
Week 20 after	-0.108*** (0.004)	-1.436*** (0.086)	-0.102*** (0.013)	-1.994*** (0.261)	-0.152*** (0.007)	-2.138*** (0.144)	-0.077*** (0.005)	-0.902*** (0.105)
Pre-ticket mean	0.27	44.858	0.279	45.753	0.362	46.951	0.199	43.244
Observations	626,430	626,430	65,606	65,606	262,282	262,282	361,352	361,352
No. of cars	16,407	16,407	2,107	2,107	13,769	13,769	14,104	14,104
R2	0.233	0.243	0.241	0.267	0.273	0.293	0.233	0.246

Notes: Regressions include car fixed effects, zone-fixed effects, and zone-specific dummy variables indicating the hour of the day, day of the week, month of the year, weekend, school holidays. They also include measures of traffic intensity and weather variables. Standard errors are two-way clustered, by car and by zone-hour.

Table A.12: Event analysis: long-run effects, analysis by the order of rides

	(1)	(2)		(3)	(4)
	<i>Long-run effects</i>			<i>Order of the ride</i>	
	Speeding	Speed		Speeding	Speed
			Ride -61 to -65	-0.018*** (0.003)	-0.26*** (0.07)
			Ride -56 to 60	-0.020*** (0.003)	-0.43*** (0.07)
			Ride -50 to -55	-0.017*** (0.003)	-0.36*** (0.07)
			Ride -46 to -50	-0.020*** (0.003)	-0.40*** (0.07)
			Ride -41 to -45	-0.020*** (0.003)	-0.34*** (0.06)
			Ride -36 to -40	-0.015*** (0.003)	-0.30*** (0.06)
			Ride -31 to -35	-0.008*** (0.003)	-0.27*** (0.06)
Month -5 before	-0.017*** (0.004)	-0.455*** (0.083)	Ride -26 to -30	-0.011*** (0.003)	-0.21*** (0.06)
Month -4 before	-0.008** (0.004)	-0.262*** (0.077)	Ride -21 to -25	-0.010*** (0.003)	-0.23*** (0.05)
Month -3 before	-0.007** (0.003)	-0.289*** (0.074)	Ride -16 to -20	-0.007*** (0.003)	-0.13*** (0.05)
Month -2 before	-0.003 (0.003)	-0.099 (0.071)	Ride -11 to -15	-0.001 (0.002)	-0.04 (0.05)
Month -1 before	-0.001 (0.003)	-0.157** (0.067)	Ride -6 to -10	-0.000 (0.002)	0.02 (0.04)
Month 1 after	-0.098*** (0.003)	-1.200*** (0.070)	Ride 1 to 5	-0.093*** (0.002)	-1.16*** (0.04)
Month 2 after	-0.099*** (0.004)	-1.251*** (0.072)	Ride 6 to 10	-0.105*** (0.002)	-1.31*** (0.04)
Month 3 after	-0.100*** (0.004)	-1.331*** (0.074)	Ride 11 to 15	-0.110*** (0.002)	-1.34*** (0.05)
Month 4 after	-0.105*** (0.004)	-1.383*** (0.071)	Ride 16 to 20	-0.115*** (0.002)	-1.47*** (0.05)
Month 5 after	-0.107*** (0.003)	-1.491*** (0.073)	Ride 21 to 25	-0.117*** (0.003)	-1.49*** (0.05)
Month 6 after	-0.103*** (0.004)	-1.327*** (0.075)	Ride 26 to 30	-0.120*** (0.003)	-1.55*** (0.05)
Month 7 after	-0.110*** (0.004)	-1.372*** (0.076)	Ride 31 to 35	-0.120*** (0.003)	-1.52*** (0.06)
Month 8 after	-0.111*** (0.004)	-1.298*** (0.078)	Ride 36 to 40	-0.123*** (0.003)	-1.56*** (0.06)
Month 9 after	-0.112*** (0.004)	-1.324*** (0.077)	Ride 41 to 45	-0.123*** (0.003)	-1.55*** (0.06)
Month 10 after	-0.114*** (0.004)	-1.377*** (0.078)	Ride 46 to 50	-0.123*** (0.003)	-1.58*** (0.06)
Month 11 after	-0.116*** (0.004)	-1.435*** (0.077)	Ride 51 to 55	-0.121*** (0.003)	-1.55*** (0.06)
Month 12 after	-0.123*** (0.004)	-1.495*** (0.076)	Ride 56 to 60	-0.128*** (0.003)	-1.61*** (0.06)
Month 13 after	-0.126*** (0.004)	-1.628*** (0.080)	Ride 61 to 65	-0.126*** (0.003)	-1.57*** (0.06)
Month 14 after	-0.128*** (0.004)	-1.726*** (0.079)	Ride 66 to 70	-0.126*** (0.003)	-1.59*** (0.06)
Month 15 after	-0.135*** (0.004)	-1.707*** (0.080)	Ride 71 to 75	-0.125*** (0.003)	-1.52*** (0.06)
Month 16 after	-0.132*** (0.004)	-1.684*** (0.078)	Ride 76 to 80	-0.127*** (0.003)	-1.68*** (0.06)
Month 17 after	-0.134*** (0.004)	-1.723*** (0.081)	Ride 81 to 85	-0.131*** (0.003)	-1.60*** (0.07)
Month 18 after	-0.133*** (0.004)	-1.677*** (0.081)	Ride 86 to 90	-0.129*** (0.003)	-1.61*** (0.07)
Month 19 after	-0.136*** (0.004)	-1.713*** (0.082)	Ride 91 to 95	-0.132*** (0.003)	-1.64*** (0.07)
Month 20 after	-0.137*** (0.004)	-1.795*** (0.087)	Ride 96 to 100	-0.129*** (0.003)	-1.69*** (0.07)
Month 21 after	-0.137*** (0.004)	-1.879*** (0.092)			
Month 22 after	-0.135*** (0.004)	-1.778*** (0.095)			
Month 23 after	-0.135*** (0.004)	-1.751*** (0.092)			
Month 24 after	-0.139*** (0.004)	-1.869*** (0.096)			
Pre-ticket mean	0.257	44.353		0.299	-3.665
Observations	991,333	991,333		1,171,931	1,171,931
No. of cars	4291	4291		16,414	16,414
R2	0.188	0.204		0.213	0.22

Notes: Regressions include car fixed effects, zone-fixed effects, and zone-specific dummy variables indicating the hour of the day, day of the week, month of the year, weekend, school holidays. They also include measures of traffic intensity and weather variables. Standard errors are two-way clustered, by car and by zone-hour.

Table A.13: Event analysis: heterogeneity I

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	<i>Car Owner</i>				<i>Payment of Ticket</i>			
	private	corporation	private	corporation	unpaid	paid	unpaid	paid
	Speeding		Speed		Speeding		Speed	
Week -11 before	0.002 (0.006)	0.002 (0.006)	0.12 (0.12)	-0.05 (0.13)	0.026** (0.011)	-0.002 (0.005)	0.31 (0.22)	-0.01 (0.09)
Week -10 before	0.006 (0.006)	-0.005 (0.006)	0.12 (0.12)	-0.28** (0.13)	0.020* (0.010)	-0.003 (0.005)	0.16 (0.24)	-0.12 (0.10)
Week -9 before	0.013** (0.006)	-0.006 (0.006)	0.07 (0.12)	-0.16 (0.13)	0.027** (0.011)	-0.001 (0.005)	0.05 (0.23)	-0.06 (0.09)
Week -8 before	0.007 (0.006)	-0.007 (0.006)	-0.03 (0.12)	-0.18 (0.13)	0.006 (0.010)	-0.001 (0.005)	-0.03 (0.22)	-0.11 (0.09)
Week -7 before	0.004 (0.006)	-0.013** (0.006)	0.05 (0.12)	-0.37*** (0.12)	0.016 (0.011)	-0.009* (0.004)	0.38* (0.21)	-0.26*** (0.09)
Week -6 before	0.010* (0.006)	-0.005 (0.006)	-0.06 (0.11)	-0.1 (0.12)	0.013 (0.010)	0.001 (0.004)	0.27 (0.21)	-0.14 (0.09)
Week -5 before	0.009 (0.006)	-0.007 (0.006)	-0.04 (0.11)	-0.18 (0.12)	0.017* (0.010)	-0.002 (0.004)	0.1 (0.21)	-0.14 (0.09)
Week -4 before	0.013** (0.006)	0.008 (0.006)	0.09 (0.11)	-0.16 (0.12)	0.030*** (0.010)	0.006 (0.004)	0.42** (0.21)	-0.12 (0.09)
Week -3 before	0.009 (0.006)	0.006 (0.006)	0.05 (0.11)	-0.01 (0.11)	0.021** (0.010)	0.005 (0.004)	0.43** (0.20)	-0.05 (0.09)
Week -2 before	0.009* (0.005)	0.005 (0.005)	0.12 (0.11)	-0.01 (0.11)	0.022** (0.009)	0.004 (0.004)	0.27 (0.20)	0.02 (0.08)
Week -1 before	0.014*** (0.005)	0.001 (0.005)	0.26** (0.10)	0 (0.11)	0.015 (0.009)	0.006 (0.004)	0.26 (0.20)	0.11 (0.08)
Week 1 after	-0.086*** (0.005)	-0.062*** (0.005)	-1.26*** (0.10)	-0.70*** (0.11)	-0.030*** (0.009)	-0.082*** (0.004)	-0.33* (0.19)	-1.10*** (0.08)
Week 2 after	-0.099*** (0.005)	-0.082*** (0.005)	-1.35*** (0.10)	-0.99*** (0.11)	-0.040*** (0.009)	-0.100*** (0.004)	-0.70*** (0.20)	-1.25*** (0.08)
Week 3 after	-0.109*** (0.005)	-0.086*** (0.006)	-1.47*** (0.11)	-1.00*** (0.11)	-0.049*** (0.010)	-0.107*** (0.004)	-0.42** (0.20)	-1.39*** (0.08)
Week 4 after	-0.105*** (0.005)	-0.085*** (0.006)	-1.35*** (0.11)	-1.16*** (0.12)	-0.054*** (0.010)	-0.103*** (0.004)	-0.66*** (0.21)	-1.37*** (0.09)
Week 5 after	-0.107*** (0.005)	-0.084*** (0.006)	-1.53*** (0.11)	-1.21*** (0.12)	-0.038*** (0.010)	-0.107*** (0.004)	-0.65*** (0.20)	-1.51*** (0.09)
Week 6 after	-0.107*** (0.006)	-0.088*** (0.006)	-1.42*** (0.11)	-1.11*** (0.12)	-0.048*** (0.010)	-0.107*** (0.004)	-0.60*** (0.22)	-1.39*** (0.09)
Week 7 after	-0.106*** (0.006)	-0.087*** (0.006)	-1.43*** (0.11)	-1.19*** (0.12)	-0.051*** (0.010)	-0.106*** (0.004)	-0.96*** (0.22)	-1.37*** (0.09)
Week 8 after	-0.102*** (0.006)	-0.089*** (0.006)	-1.37*** (0.11)	-1.09*** (0.13)	-0.048*** (0.011)	-0.105*** (0.004)	-0.53** (0.22)	-1.36*** (0.09)
Week 9 after	-0.101*** (0.006)	-0.093*** (0.006)	-1.45*** (0.11)	-1.04*** (0.12)	-0.055*** (0.011)	-0.106*** (0.004)	-0.74*** (0.22)	-1.34*** (0.09)
Week 10 after	-0.107*** (0.006)	-0.094*** (0.006)	-1.47*** (0.11)	-1.24*** (0.12)	-0.063*** (0.010)	-0.107*** (0.004)	-0.82*** (0.22)	-1.46*** (0.09)
Week 11 after	-0.110*** (0.006)	-0.088*** (0.006)	-1.44*** (0.11)	-1.09*** (0.12)	-0.055*** (0.011)	-0.107*** (0.004)	-0.81*** (0.23)	-1.34*** (0.09)
Week 12 after	-0.108*** (0.006)	-0.090*** (0.006)	-1.60*** (0.12)	-1.37*** (0.12)	-0.059*** (0.011)	-0.107*** (0.004)	-1.19*** (0.24)	-1.54*** (0.09)
Week 13 after	-0.105*** (0.006)	-0.093*** (0.006)	-1.41*** (0.12)	-1.09*** (0.12)	-0.068*** (0.011)	-0.104*** (0.004)	-0.80*** (0.23)	-1.34*** (0.09)
Week 14 after	-0.108*** (0.006)	-0.097*** (0.006)	-1.47*** (0.12)	-1.11*** (0.12)	-0.065*** (0.011)	-0.109*** (0.004)	-0.91*** (0.24)	-1.36*** (0.09)
Week 15 after	-0.107*** (0.006)	-0.092*** (0.006)	-1.57*** (0.11)	-1.15*** (0.12)	-0.070*** (0.011)	-0.105*** (0.004)	-0.94*** (0.22)	-1.43*** (0.09)
Week 16 after	-0.113*** (0.006)	-0.106*** (0.006)	-1.62*** (0.11)	-1.28*** (0.12)	-0.089*** (0.011)	-0.114*** (0.005)	-0.97*** (0.22)	-1.54*** (0.09)
Week 17 after	-0.106*** (0.006)	-0.096*** (0.006)	-1.38*** (0.12)	-1.32*** (0.13)	-0.090*** (0.011)	-0.103*** (0.005)	-1.12*** (0.22)	-1.39*** (0.09)
Week 18 after	-0.114*** (0.006)	-0.103*** (0.006)	-1.53*** (0.12)	-1.23*** (0.12)	-0.080*** (0.011)	-0.115*** (0.004)	-0.94*** (0.23)	-1.46*** (0.09)
Week 19 after	-0.115*** (0.006)	-0.105*** (0.006)	-1.62*** (0.12)	-1.41*** (0.12)	-0.094*** (0.011)	-0.114*** (0.004)	-1.29*** (0.23)	-1.55*** (0.09)
Week 20 after	-0.109*** (0.006)	-0.107*** (0.006)	-1.49*** (0.12)	-1.40*** (0.13)	-0.078*** (0.012)	-0.114*** (0.004)	-1.00*** (0.23)	-1.52*** (0.09)
Pre-ticket mean	0.262	0.278	44.74	44.97	0.274	0.269	45.099	44.809
No. of cars	8,393	8,014	8,393	8,014	2,474	13,933	2,474	13,933
Observations	312,885	313,545	312,885	313,545	100,364	526,066	100,364	526,066
R2	0.237	0.232	0.25	0.24	0.243	0.231	0.25	0.24

Notes: Regressions include car fixed effects, zone-fixed effects, and zone-specific dummy variables indicating the hour of the day, day of the week, month of the year, weekend, school holidays. They also include measures of traffic intensity and weather variables. Standard errors are two-way clustered, by car and by zone-hour.

Table A.14: Event analysis: heterogeneity II

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	<i>by driving frequency</i>				<i>by number plate region</i>					
	Infrequent	Frequent	Infrequent	Frequent	Local region	Prague	Other regions	Local region	Prague	Other regions
	Speeding		Speed		Speeding			Speed		
Week -11 before	0.006 (0.013)	0.002 (0.005)	0.28 (0.24)	0.02 (0.09)	-0.006 (0.006)	0.011* (0.006)	0.001 (0.013)	-0.1 (0.13)	0.19 (0.13)	0.03 (0.26)
Week -10 before	-0.001 (0.013)	0.001 (0.004)	0.13 (0.24)	-0.09 (0.10)	-0.001 (0.006)	0.002 (0.006)	0.001 (0.012)	-0.04 (0.13)	-0.11 (0.14)	-0.10 (0.26)
Week -9 before	0.013 (0.013)	0.004 (0.004)	-0.16 (0.24)	-0.02 (0.09)	0.007 (0.006)	-0.002 (0.006)	0.015 (0.013)	0.07 (0.13)	-0.18 (0.13)	0.08 (0.26)
Week -8 before	0.001 (0.013)	0.001 (0.004)	0.05 (0.24)	-0.11 (0.09)	0 (0.006)	0.003 (0.006)	-0.008 (0.012)	-0.1 (0.13)	-0.1 (0.13)	-0.26 (0.26)
Week -7 before	0.009 (0.012)	-0.005 (0.004)	-0.26 (0.24)	-0.14 (0.09)	-0.006 (0.006)	-0.003 (0.006)	-0.004 (0.011)	-0.11 (0.13)	-0.18 (0.13)	-0.20 (0.24)
Week -6 before	0.012 (0.012)	0.002 (0.004)	0.02 (0.22)	-0.08 (0.09)	-0.003 (0.006)	0.005 (0.006)	0.015 (0.012)	-0.14 (0.13)	-0.04 (0.12)	0.00 (0.24)
Week -5 before	-0.01 (0.012)	0.003 (0.004)	-0.22 (0.22)	-0.08 (0.09)	0.005 (0.006)	-0.003 (0.006)	0.004 (0.012)	-0.17 (0.12)	-0.01 (0.12)	-0.19 (0.24)
Week -4 before	0.025** (0.011)	0.009** (0.004)	0.02 (0.22)	-0.04 (0.09)	0.007 (0.006)	0.014** (0.006)	0.01 (0.011)	0.01 (0.12)	0.01 (0.12)	-0.27 (0.24)
Week -3 before	0.008 (0.011)	0.008* (0.004)	0.21 (0.21)	0 (0.09)	0.004 (0.006)	0.009 (0.006)	0.015 (0.011)	-0.07 (0.12)	0.11 (0.12)	0.05 (0.23)
Week -2 before	0.018 (0.011)	0.006 (0.004)	0.28 (0.21)	0.03 (0.08)	0.005 (0.006)	0.007 (0.006)	0.017 (0.011)	-0.06 (0.12)	0.18 (0.11)	0.07 (0.22)
Week -1 before	0.016 (0.010)	0.007* (0.004)	0.23 (0.20)	0.11 (0.08)	0.006 (0.006)	0.004 (0.006)	0.025** (0.010)	0.12 (0.12)	0.16 (0.12)	0.06 (0.21)
Week 1 after	-0.085*** (0.010)	-0.072*** (0.004)	-1.20*** (0.20)	-0.93*** (0.08)	-0.085*** (0.006)	-0.063*** (0.006)	-0.071*** (0.011)	-1.10*** (0.11)	-0.82*** (0.12)	-1.10*** (0.22)
Week 2 after	-0.112*** (0.010)	-0.087*** (0.004)	-1.47*** (0.20)	-1.11*** (0.08)	-0.095*** (0.006)	-0.082*** (0.006)	-0.105*** (0.011)	-1.28*** (0.11)	-1.07*** (0.12)	-1.12*** (0.20)
Week 3 after	-0.123*** (0.010)	-0.093*** (0.004)	-1.48*** (0.20)	-1.19*** (0.08)	-0.106*** (0.006)	-0.083*** (0.006)	-0.113*** (0.011)	-1.42*** (0.11)	-1.00*** (0.12)	-1.42*** (0.21)
Week 4 after	-0.097*** (0.010)	-0.094*** (0.004)	-1.43*** (0.20)	-1.22*** (0.09)	-0.100*** (0.006)	-0.087*** (0.006)	-0.104*** (0.011)	-1.38*** (0.12)	-1.07*** (0.12)	-1.47*** (0.22)
Week 5 after	-0.110*** (0.010)	-0.093*** (0.004)	-1.64*** (0.21)	-1.32*** (0.09)	-0.106*** (0.006)	-0.081*** (0.006)	-0.109*** (0.011)	-1.51*** (0.12)	-1.15*** (0.12)	-1.70*** (0.22)
Week 6 after	-0.110*** (0.011)	-0.095*** (0.004)	-1.36*** (0.21)	-1.25*** (0.09)	-0.106*** (0.006)	-0.084*** (0.006)	-0.108*** (0.011)	-1.50*** (0.12)	-1.01*** (0.12)	-1.29*** (0.22)
Week 7 after	-0.108*** (0.011)	-0.095*** (0.004)	-1.57*** (0.21)	-1.26*** (0.09)	-0.106*** (0.006)	-0.086*** (0.006)	-0.098*** (0.012)	-1.39*** (0.12)	-1.14*** (0.12)	-1.54*** (0.23)
Week 8 after	-0.110*** (0.011)	-0.093*** (0.004)	-1.37*** (0.20)	-1.20*** (0.09)	-0.102*** (0.006)	-0.083*** (0.006)	-0.120*** (0.012)	-1.34*** (0.13)	-0.96*** (0.13)	-1.77*** (0.24)
Week 9 after	-0.111*** (0.011)	-0.095*** (0.004)	-1.40*** (0.21)	-1.22*** (0.09)	-0.105*** (0.006)	-0.082*** (0.006)	-0.128*** (0.012)	-1.48*** (0.12)	-0.86*** (0.12)	-1.76*** (0.23)
Week 10 after	-0.123*** (0.011)	-0.097*** (0.004)	-1.75*** (0.22)	-1.30*** (0.09)	-0.104*** (0.006)	-0.091*** (0.006)	-0.122*** (0.013)	-1.43*** (0.12)	-1.17*** (0.12)	-1.77*** (0.26)
Week 11 after	-0.109*** (0.011)	-0.097*** (0.004)	-1.32*** (0.21)	-1.25*** (0.09)	-0.107*** (0.006)	-0.086*** (0.006)	-0.115*** (0.012)	-1.35*** (0.12)	-1.12*** (0.13)	-1.47*** (0.24)
Week 12 after	-0.112*** (0.011)	-0.097*** (0.004)	-1.55*** (0.22)	-1.47*** (0.09)	-0.108*** (0.006)	-0.084*** (0.006)	-0.122*** (0.012)	-1.55*** (0.13)	-1.38*** (0.13)	-1.57*** (0.25)
Week 13 after	-0.127*** (0.011)	-0.094*** (0.005)	-1.75*** (0.21)	-1.17*** (0.09)	-0.104*** (0.006)	-0.086*** (0.006)	-0.122*** (0.012)	-1.32*** (0.13)	-1.15*** (0.13)	-1.31*** (0.24)
Week 14 after	-0.116*** (0.011)	-0.100*** (0.004)	-1.62*** (0.21)	-1.23*** (0.09)	-0.113*** (0.006)	-0.088*** (0.006)	-0.113*** (0.012)	-1.43*** (0.13)	-1.06*** (0.13)	-1.49*** (0.23)
Week 15 after	-0.113*** (0.011)	-0.097*** (0.005)	-1.60*** (0.22)	-1.31*** (0.09)	-0.109*** (0.006)	-0.088*** (0.006)	-0.106*** (0.012)	-1.45*** (0.12)	-1.20*** (0.13)	-1.53*** (0.22)
Week 16 after	-0.121*** (0.011)	-0.107*** (0.005)	-1.65*** (0.21)	-1.41*** (0.09)	-0.113*** (0.006)	-0.102*** (0.006)	-0.120*** (0.012)	-1.61*** (0.12)	-1.22*** (0.13)	-1.67*** (0.24)
Week 17 after	-0.112*** (0.011)	-0.099*** (0.005)	-1.74*** (0.22)	-1.28*** (0.09)	-0.105*** (0.006)	-0.092*** (0.007)	-0.116*** (0.012)	-1.41*** (0.13)	-1.22*** (0.13)	-1.60*** (0.24)
Week 18 after	-0.124*** (0.011)	-0.106*** (0.004)	-1.73*** (0.22)	-1.31*** (0.09)	-0.113*** (0.006)	-0.098*** (0.006)	-0.128*** (0.012)	-1.36*** (0.13)	-1.31*** (0.13)	-1.73*** (0.26)
Week 19 after	-0.137*** (0.011)	-0.106*** (0.004)	-1.85*** (0.22)	-1.46*** (0.09)	-0.113*** (0.006)	-0.103*** (0.006)	-0.128*** (0.012)	-1.50*** (0.13)	-1.39*** (0.13)	-2.06*** (0.24)
Week 20 after	-0.116*** (0.011)	-0.106*** (0.005)	-1.77*** (0.22)	-1.39*** (0.09)	-0.113*** (0.006)	-0.095*** (0.007)	-0.133*** (0.012)	-1.53*** (0.13)	-1.20*** (0.13)	-1.96*** (0.25)
Pre-ticket mean	0.321	0.261	46.48	44.57	0.258	0.273	0.299	44.48	44.96	45.74
Observations	88,557	537,873	88,557	537,873	278,333	276,598	71,499	278,333	276,598	71,499
No. of cars	8,148	8,259	8,148	8,259	5,860	7,817	2,730	5,860	7,817	2,730
R2	0.306	0.219	0.31	0.23	0.223	0.237	0.26	0.23	0.25	0.28

Notes: Regressions include car fixed effects, zone-fixed effects, and zone-specific dummy variables indicating the hour of the day, day of the week, month of the year, weekend, school holidays. They also include measures of traffic intensity and weather variables. Standard errors are two-way clustered, by car and by zone-hour.

Table A.15: Event analysis: reoffending status and second ticket

	(1)	(2)	(3)	(4)	(5)	(6)
	<i>First Ticket</i>				<i>Second Ticket</i>	
	Did not reoffend		Reoffended		Reoffended	
	Speeding	Speed	Speeding	Speed	Speeding	Speed
Week -11 before	0.009* (0.005)	0.21** (0.10)	0.004 (0.009)	0.15 (0.18)	0.007 (0.010)	-0.01 (0.20)
Week -10 before	0.008 (0.005)	0.10 (0.10)	0.001 (0.009)	0.14 (0.20)	-0.004 (0.011)	-0.23 (0.20)
Week -9 before	0.011** (0.005)	0.15 (0.10)	-0.000 (0.009)	-0.03 (0.19)	-0.006 (0.011)	-0.12 (0.20)
Week -8 before	0.007 (0.005)	0.14 (0.10)	0.003 (0.009)	-0.03 (0.18)	0.008 (0.010)	-0.13 (0.21)
Week -7 before	0.001 (0.005)	0.06 (0.10)	0.007 (0.009)	0.07 (0.20)	0.001 (0.010)	0.05 (0.20)
Week -6 before	0.009* (0.005)	0.07 (0.10)	0.018** (0.009)	0.26 (0.19)	-0.004 (0.010)	-0.22 (0.20)
Week -5 before	0.004 (0.005)	0.08 (0.10)	0.011 (0.009)	-0.05 (0.18)	0.005 (0.010)	0.07 (0.19)
Week -4 before	0.016*** (0.005)	0.10 (0.10)	0.014 (0.009)	0.06 (0.17)	-0.004 (0.009)	-0.00 (0.19)
Week -3 before	0.009** (0.005)	0.14 (0.09)	0.024*** (0.009)	0.24 (0.17)	-0.002 (0.009)	0.17 (0.19)
Week -2 before	0.011** (0.005)	0.20** (0.09)	0.010 (0.008)	0.06 (0.17)	-0.006 (0.010)	-0.06 (0.19)
Week -1 before	0.012*** (0.004)	0.19** (0.09)	0.007 (0.008)	0.21 (0.17)	-0.009 (0.010)	-0.35* (0.20)
Week 1 after	-0.076*** (0.004)	-0.94*** (0.09)	-0.031*** (0.008)	-0.08 (0.17)	-0.059*** (0.009)	-0.85*** (0.19)
Week 2 after	-0.095*** (0.005)	-1.04*** (0.09)	-0.038*** (0.008)	-0.29* (0.17)	-0.072*** (0.010)	-1.03*** (0.20)
Week 3 after	-0.101*** (0.005)	-1.18*** (0.09)	-0.046*** (0.009)	-0.30* (0.18)	-0.072*** (0.010)	-1.15*** (0.20)
Week 4 after	-0.095*** (0.005)	-1.12*** (0.09)	-0.052*** (0.008)	-0.47*** (0.17)	-0.081*** (0.010)	-1.13*** (0.20)
Week 5 after	-0.098*** (0.005)	-1.22*** (0.09)	-0.052*** (0.009)	-0.63*** (0.17)	-0.068*** (0.011)	-0.95*** (0.20)
Week 6 after	-0.096*** (0.005)	-1.16*** (0.10)	-0.054*** (0.009)	-0.38** (0.18)	-0.083*** (0.010)	-1.38*** (0.22)
Week 7 after	-0.094*** (0.005)	-1.24*** (0.10)	-0.054*** (0.009)	-0.41** (0.17)	-0.081*** (0.010)	-1.31*** (0.21)
Week 8 after	-0.090*** (0.005)	-1.09*** (0.10)	-0.059*** (0.009)	-0.41** (0.18)	-0.080*** (0.010)	-1.20*** (0.22)
Week 9 after	-0.093*** (0.005)	-1.15*** (0.10)	-0.060*** (0.008)	-0.32* (0.17)	-0.061*** (0.010)	-0.61*** (0.21)
Week 10 after	-0.100*** (0.005)	-1.24*** (0.10)	-0.054*** (0.008)	-0.59*** (0.18)	-0.065*** (0.010)	-0.93*** (0.20)
Week 11 after	-0.095*** (0.005)	-1.13*** (0.10)	-0.058*** (0.009)	-0.53*** (0.18)	-0.074*** (0.010)	-1.09*** (0.22)
Week 12 after	-0.094*** (0.005)	-1.26*** (0.10)	-0.060*** (0.008)	-0.78*** (0.18)	-0.068*** (0.010)	-0.96*** (0.21)
Week 13 after	-0.096*** (0.005)	-1.09*** (0.10)	-0.056*** (0.009)	-0.50*** (0.18)	-0.078*** (0.010)	-1.29*** (0.21)
Week 14 after	-0.098*** (0.005)	-1.15*** (0.10)	-0.059*** (0.009)	-0.39** (0.19)	-0.080*** (0.011)	-1.29*** (0.20)
Week 15 after	-0.093*** (0.005)	-1.15*** (0.10)	-0.060*** (0.009)	-0.63*** (0.17)	-0.088*** (0.010)	-1.49*** (0.22)
Week 16 after	-0.104*** (0.005)	-1.24*** (0.10)	-0.074*** (0.009)	-0.90*** (0.18)	-0.074*** (0.011)	-1.27*** (0.22)
Week 17 after	-0.098*** (0.005)	-1.20*** (0.10)	-0.054*** (0.009)	-0.57*** (0.18)	-0.081*** (0.011)	-1.32*** (0.22)
Week 18 after	-0.102*** (0.005)	-1.15*** (0.10)	-0.072*** (0.008)	-0.78*** (0.18)	-0.082*** (0.010)	-1.27*** (0.22)
Week 19 after	-0.105*** (0.005)	-1.36*** (0.10)	-0.070*** (0.009)	-0.85*** (0.18)	-0.083*** (0.010)	-1.33*** (0.21)
Week 20 after	-0.100*** (0.005)	-1.25*** (0.10)	-0.075*** (0.009)	-0.78*** (0.17)	-0.071*** (0.011)	-1.25*** (0.21)
Pre-ticket mean	0.245	44.412	0.283	44.716	0.277	45.117
No. of cars	12,802	12,802	2,551	2,551	1,694	1,694
Observations	417,829	417,829	143,292	143,292	101,530	101,530
R2	0.228	0.24	0.236	0.24	0.227	0.23

Notes: Regressions include car fixed effects, zone-fixed effects, and zone-specific dummy variables indicating the hour of the day, day of the week, month of the year, weekend, school holidays. They also include measures of traffic intensity and weather variables. Standard errors are two-way clustered, by car and by zone-hour.

B Appendix

B.1 Mapping experiences in expectations

Making use of (3), we can substitute for $p^{t-1}(s)$ and obtain

$$p^t(s) = P(\{s_{t-1}, \mathbb{T}^t(s_{t-1})\}, \{s_{t-2}, \mathbb{T}^t(s_{t-2})\}, \dots, \{s_0, \mathbb{T}^t(s_0)\}, \\ P(\{s_{t-2}, \mathbb{T}^{t-1}(s_{t-2})\}, \dots, \{s_0, \mathbb{T}^{t-1}(s_0)\}, p^{t-2}(s))) = \dots \quad (\text{B.1})$$

Iterating this substitution and accounting for the fact that a ride from period τ could, in principle, result in a ticket that is delivered in any period $t > \tau$, we arrive at

$$= \Pi_t(\{s_{t-1}, \mathbb{T}^t(s_{t-1})\}, \{s_{t-2}, \mathbb{T}^t(s_{t-2}), \mathbb{T}^{t-1}(s_{t-2})\}, \dots, \\ \{s_0, \mathbb{T}^t(s_0), \mathbb{T}^{t-1}(s_0), \dots, \mathbb{T}^1(s_0)\}, p^0(s)) \quad (\text{B.2})$$

Let us define the vector $\vec{\mathbb{T}}(t, s_\tau) := (\mathbb{T}^t(s_\tau), \mathbb{T}^{t-1}(s_\tau), \dots, \mathbb{T}^{\tau+1}(s_\tau))$, which captures a sequence of ‘ticketing experiences’ (i.e. receiving or not receiving a ticket) that follow from a ride in period $\tau < t$ at speed s_τ which might result in a ticket arriving in any period $\tau + 1, \tau + 2, \dots, t - 1, t$. With this notation, we arrive at

$$p^t(s) = \Pi_t\left(\left(\{s_{t-1}, \vec{\mathbb{T}}(t, s_{t-1})\}\right), \dots, \left(\{s_1, \vec{\mathbb{T}}(t, s_1)\}\right), \left(\{s_0, \vec{\mathbb{T}}(t, s_0)\}\right), p^0(s)\right), \quad (\text{B.3})$$

which is the mapping from (4).


B.2 Assignment and Outcome Period in the RDD

Figure B.1 illustrates our approach (introduced in Section 4.1) to defining an assignment speed and an outcome period for each car i . The figure depicts the driving pattern and speed of two cars. Recall first that the a -month long *assignment period* (with, e.g., $a = 4$ months), begins the first time a car is observed. For both cars in this example we observe the highest speed during the assignment period on the same day (a little more than 2 months after their first ride). This point in time – the trigger day – is indicated by the vertical dashed line. The two cars’ *maximum speed* during the assignment period, S_i and S_j , define their *assignment speed*. In the Figure, the one car (indicated with \times) has an assignment speed above the enforcement cutoff and will thus trigger a speeding ticket. For the other car (indicated with \bullet), the assignment speed S_j is below the enforcement cutoff (indicated with the dashed, horizontal line at 14km/h above the speed limit). The latter car will not receive a ticket (unless it reaches a speed above the enforcement limit at a later point in time (after the end of the assignment period)).

Based on all observations (not only the two cars depicted in the Figure), we would now compute the earliest day a ticket for a speeding offense from this trigger day is sent. This ‘shortest delay’ (in the example: about two weeks) then defines the start of an f -month *outcome period* (e.g., $f = 4$). For all cars that have an *assignment speed* recorded on the same trigger day, we would define an

identical *outcome period*. Importantly, this property holds independent of the level of the cars' assignment speed (in particular, independently of whether S_i or S_j is below or above the first or the second cutoff from the RDD).

Figure B.1: Illustration of Assignment and Outcome Periods



Notes: The figure illustrates our approach to defining assignment and outcome periods for the driving patterns of two cars.

The Figure above also illustrates that the recorded speed within the assignment period but *after* the trigger date must be lower by definition. To the extent that the end of the assignment period overlaps with the outcome period – which is more likely to occur if the trigger is observed early during the assignment period (and/or if the ticket delay is rather short) – this will result in lower speed values during the (early phase of the) outcome period. Note, however, that this property once again holds symmetrically for cars with assignment speeds below or above any of the two cutoffs. Hence, this property does not drive our RDD estimates. (The latter point can be easily verified in robustness checks. When we define alternative outcome windows that do not overlap with the assignment period, we obtain essentially the same results.)

Let us finally note that our approach defines assignment and outcome periods in a flexible, car-specific way (i.e. relative to the first observation and relative to the trigger day). In an earlier version of this paper we adopted a more static strategy that simply defined the initial months of the sample as the assignment window and the latter months as the outcome period. This static approach produced very similar results, but explored a much smaller part of the sample.

B.3 Mean reversion in event analysis

To illustrate the mean reversion issue in the raw data (which is captured in Figure 7), we introduce a simple framework of speed choices. Speed s_{it} on a drive-through at time t by car i is given by

$$s_{it} = \lambda_i + \beta T_{it} + \gamma X_{it} + \varepsilon_{it}, \quad (\text{B.4})$$

where λ_i is a car fixed effect, T_{it} is a treatment dummy (indicating that the driver has received a ticket prior to time t), X_{it} is a vector of exogenous variables and ε_{it} is an error term. By definition, the driver had to commit a speeding violation in order to receive a ticket. Hence, the ticket dummy is only positive if the car was driving above the enforcement cutoff k^1 at some point t' . This can be written as

$$T_{it} = 1 \iff y_{it} > k^1 \text{ for some } t' < 0, \quad (\text{B.5})$$

where $t = 0$ indicates the event (ticket delivery). It thus follows that

$$T_{it} = 1 \iff \varepsilon_{it} > k^1 - (\lambda_i + \beta T_{it} + \gamma X_{it}) \text{ for some } t' < 0 \quad (\text{B.6})$$

By construction, the trigger observation with an unusually high draw of ε_{it} occurs among the pre-ticket observations. In practice, most trigger observations are concentrated during the three weeks prior to receiving the ticket, as most tickets were received within three weeks after the speeding offense. As noted in the main text, this leads to the increase in speeding occurrences observed in the raw data. In contrast, the observations after the ticket have, by assumption, ε_{it} drawn from a mean-zero distribution. This implies that there is a negative correlation between the treatment dummy and the error term. Neglecting this issue would result in an overestimation of the speeding tickets' effect.