Learning from Law Enforcement*

Libor Dušek† and Christian Traxler‡

This version: August 5, 2019

Comments welcome!

Abstract

This paper studies how punishment for past offenses shapes future compliance behavior via learning. The context of our study is traffic law enforcement through automated speed cameras. We use unique data on speeding tickets and full driving histories of more than one million cars tracked over several years in a suburb of Prague. In our setting, punishment neither implies incapacitation nor do past tickets alter the 'price' for future offenses. This allows us to identify specific deterrence effects induced by learning from law enforcement. We present results from two empirical strategies. Firstly, a regression discontinuity design exploits two speed level cutoffs which provide variation in punishment at the extensive (receiving a speeding ticket) and intensive margin (tickets with low or high fines), respectively. The RDD reveals strong and precisely estimated responses to speeding tickets: the speeding rate drops by a third (10 percentage points) and chances of getting a further ticket fall by 70%. An increase in punishment at the intensive margin – a more than a doubling of fines – triggers only a limited additional effect. Secondly, an event study makes use of the high-frequency nature of our data. The average treatment effects on the treated obtained from the event study confirms all LATEs from the RDD. We also document that driving responses are immediate and very persistent over time. Even two years after receiving a ticket there is no evidence on 'backsliding' towards speeding. The results reject unlearning and temporary salience effects and support a reinforcement learning model in which agents, after experiencing punishment, update their priors on the expected costs of future offending in a discontinuous, 'coarse' manner. Additional results indicate that learning from (local) 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 and Michael Mueller-Smith. We would also like to thank Isabel Anaya, Jiri Krejsa, Marcel Tkacik, Jan Vavra and Paulina Ockajova for excellent research assistance. The constructive cooperation with the town hall officials of Ricany as well as the research funding by the GACR (grant 17-16583J) and the DFG (grant TR 1471/1-1) is highly appreciated.

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

‡Hertie School of Governance, Berlin, and CESifo. traxler@hertie-school.org.

1 Introduction

How do individuals respond to being punished for a crime or an offense? There are typically numerous mechanisms through which most forms of punishment can shape future behavior. Imprisonment, for instance, can imply incapacitation and peer effects. If one excludes such implicit mechanisms and considers a context where the punishment for future delinquencies is constant (rather than increasing in, e.g., earlier convictions), 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 priors about relevant parameters of the enforcement process (e.g., the detection risk, the level of punishment). The agent learns from law enforcement and, in turn, adjusts behavior. A learning framework therefore 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. The isolation of a learning channel, however, is typically impossible. As pointed out above, (expected) punishment is often increasing with prior convictions. Past punishment then implies a general deterrence effect. More importantly, punishment typically comprises a multidimensional treatment that influences later behavior along many different channels. Imprisonment may imply incapacitation and aging (Ganong, 2012; Barbarino and Mastrobuoni, 2014), criminogenic peer effects (Bayer et al., 2009), or diverging labor market consequences (Mueller-Smith, 2015; Bhuller et al., 2019). Monetary forms of punishment, in contrast, may involve non-trivial income effects. Large fines imposed on, e.g., detected tax evaders could then alter future behavior by influencing their risk tolerance (Kolm, 1973).

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 allow us to track the driving histories of more than 1.2 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 observe, e.g., only rearresting. In our context there are no relevant income effects and the price for re-offending is unaffected by past punishment: fines are modest (around \$40–85) and independent of past tickets. Insurance rates remain constant, too. In addition, punishment does not induce incapacitation (driving licenses are never revoked or suspended). Thus the set-up is ideal to isolate if and how agents learn from (traffic) law enforcement.

[1]The 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.

We provide complementary results from two research designs. First, we implement a regression discontinuity design (RDD) which exploits two speed-level cutoffs. The first cutoff is an enforcement threshold used by the local police forces. If a car's speed is above this cutoff (14km/h above the speed limit), the system automatically triggers a speeding ticket. If the speed is above a second threshold, the fine for the speeding ticket more than doubles (from \$40 to \$85). The two cutoffs thus offer variation in punishment at the extensive (receiving or not receiving a speeding ticket) and intensive margin (tickets with low or high fines), respectively. Based on the latter variation, we can study whether specific deterrence effects are increasing with higher fines.

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%); the chances of getting a (further) ticket decline by 70%. These numbers reflect a pronounced shift in the speed distribution. The mass of rides above the speed limit (also 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%, with larger changes at the top of the speed distribution.

These findings are consistent with a reinforcement learning framework. After receiving a speeding ticket, drivers update their priors on $q(s)$, the expected penalties for driving at speed s, and accordingly adjust their optimal speed. Within this framework, one can distinguish between 'fine-grained' and more 'coarse' updating of priors. For the former case, we predict nuanced speed adjustments and, eventually, bunching below the enforcement cutoff. Under coarse updating, $q(s)$ raises for almost any speed above the limit. In turn, we should observe a strong drop in speeding (with heaping at the speed limit) and too little experimentation to figure out the enforcement cutoff. The evidence is consistent with these latter predictions: drivers learn about the enforcement of speed limits and it's consequences – but not about the threshold used in the enforcement of speeding violations.

Coarse updating (in contrast to the fine-grained case) further implies that higher fines do *not*, on average, amplify the specific deterrence effects of speeding tickets. In line with this prediction, we do not find statically significant additional effects at the second cutoff (where the fine increases from \$40 to \$85). Only for a theory-motivated subsample – rides observed under driving conditions that favour speeding – we detect stronger effects of tickets with higher fines. Overall, the evidence suggests that the variation in fines has a minor influence on the way drivers learn their lesson. Intensive margin variation in punishment has therefore only a limited effect.

We complement the RDD results with an event study which makes use of the high-frequency resolution of our data. Based on information on the exact timing when a speeding ticket is delivered, we explore within-car variation before and after receiving a ticket. The results, firstly, corroborate all findings from the RDD. In fact, the average treatment effects on the treated (ATTs) obtained from the event study are almost identical to the local average treatment effects (LATEs) from the RDD. Secondly, we document 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. These results contradict 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). The evidence clearly rejects such unlearning. The responses to punishment are persistent and not a temporary salience effect, which could emerge in models with limited attention or cognition (Gabaix, 2019).

In an additional step, we study how broadly (or narrowly) drivers adjust their behavior. We document that a drop in speeding rates is also observed for other speed camera zones – beyond the one that triggered the ticket. Studying the time spent on trips on the un-monitored stretch between two speed cameras, we provide further evidence suggesting that the increased compliance with speed limits inside speed camera zones is not compensated by more aggressive speeding outside these monitored zones. Experiencing punishment seems to trigger a relatively broad behavioral adjustment.

Finally, we also identify spillover effects. Speeding tickets not only induce a slow down of treated cars but also of those that travel in lines behind 'ticketed' cars. Depending on traffic density (and the composition of lines), this spillover can reach well beyond the next car behind a treated one. In addition to these (at least partially mechanical) 'backward spillovers', we also find some evidence on 'forward spillovers'. Speeding tickets make (otherwise aggressive) cars drive slower; driving less pushy, in turn, results in the car traveling ahead of a treated one at a slower pace, too. The latter effect, however, is weaker and less robust.

Our paper relates to several strands of research. First of all, we contribute to the literature on learning (Mobius and Rosenblat, 2014). 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 (e.g., Lochner, 2007).[2] While there is also some evidence indicating that word-of-mouth learning *between* peers influences compliance decisions (Rincke and Traxler, 2011; Drago et al., 2019), we provide evidence that identifies responses to punishment mediated by *within*-agent learning.[3]

In doing so, we also contribute to the economic analysis of deterrence (Chalfin and McCrary, 2017). General deterrence effects are theoretically well understood 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, Section 7) 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,

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

[3] For the domain of tax enforcement, a recent survey documents strong correlations between past exposure to tax audits and firms perceived auditing risk (Bérgolo et al., 2018). This is consistent with randomized audit interventions which typically induce long-lasting compliance responses among audited taxpayers (see, e.g., Kleven et al. 2011; Advani et al. 2019, who find positive effects, but also DeBacker et al. 2015, who report negative effects). The repeated interaction of tax authorities and taxpayers, however, is fairly complex. Hence, in addition to the income effects mentioned see above, there are numerous mechanism beyond learning and updated expectations that could shape responses to experiencing an audit.

2007; Drago et al., 2011; Ganong, 2012; Kuziemko, 2013; Mastrobuoni and Terlizzese, 2019).[4] 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.[5]

We differ from (and contribute to) this strand of research in several ways. First, our set-up – the enforcement of speed limits – 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 precisely track behavioral responses to punishment over time. 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).[6] Based on this feature, we illustrate changes in a continuous distribution of (non-)compliance behavior.

Third, using a 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 obtain limited and less precisely estimated effects. Fourth, we present a simple, formal framework which offers a coherent interpretation of these responses (to both extensive and intensive margin variation in punishment) in terms of reinforcement learning.[7] The analysis shows that 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 (with Hansen, 2015, further discussed below, as an important exception). Note that studies on prison sentences compare imprisonment with other, less severe forms of punishment (or different imprisonment conditions) – thus focusing mainly on the intensive margin. Exogenous extensive margin variation in punishment, as it is explored in our context, is thus rare.

We are convinced that the relevance of learning-induced deterrence extends to other domains. Support for this view is provided by Philippe (2019), who studies a recent reform of minimum sentencing requirements in France. The requirements only applied to a specific form of recidivism: committing exactly the same type of crime again. Using a difference-in-difference strategy that compares offenders who faced the same sentencing threat, Philippe (2019) finds that only those who had the chance to learn about the exact scope of the reform, respond strategically, by

[4] See Nagin et al. (2009) for a comprehensive review of this research.

[5] Mixed results are also reported in the literature on tax enforcement (see fn. 3).

[6] This feature also distinguishes our work from other research on traffic violations, discussed below.

[7] The framework complements earlier analyses, which mainly focussed on the social dimension of learning (Sah, 1991) or on Bayesian learning (Lochner, 2001). Our analysis further differs from these studies in that we examine a continuous rather than a binary (non-)compliance decision.

committing fewer crimes of the same type. Correlational evidence, which compares offenders who did or did not show-up in court (and did not get the law explained), further supports this interpretation. The chances to learn might therefore shape individual responses to interactions with the criminal justice system.

Finally, our study also adds to the economic research on traffic law enforcement. In this context, a few quasi-experimental studies estimate effects on the extensive margin of 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, however, in that past punishment carries a general deterrent effect (related to increases in future penalties).

Hansen (2015) also studies 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 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. Two features distinguish our paper from all these studies: the simplicity of our set-up (which neither includes suspensions nor jail sentences, etc.) 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 distributional effects 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, with no evidence on compensatory speeding on un-monitored parts of the road. The spillover effects further imply that the tickets' impact spreads to a much vaster population of cars. Together with a potential general deterrence effect of the speed cameras, this contributes to an overall decline in travel speed. These findings appear relevant, given that the WHO (2018) considers effective speed management policies as the central strategy to reduce 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; Tang, 2017) indicate that the observed decline in speeding might not only lower accident risks but 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.

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

Second, and more broadly, our work points to the importance of learning and information policies in mediating deterrence effects. In a similar vein, Philippe (2019) notes that differences in informational policies might actually explain why similar policy changes had different effects, depending on whether or not the reforms were largely publicized (or individually communicated). (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 results also highlight the potential value of ambiguity. In our context, the exact speed cutoff above which speeding triggers a ticket (14km/h above the speed limit) is unknown. The 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 may therefore amplify behavioral responses. In principle, this result might apply to numerous domains – such as petty theft, minor drug possession, public nuisances, tax evasion, or environmental pollution – where offenders face ambiguity about the exact point up to which authorities 'tolerate' illegal behavior and where punishment starts.[9]

The remainder of the paper is structured as follows. After describing the institutional background and our data, Section 3 introduces a 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 of 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 at 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 a given zone record a car's number plate together with a precise time stamp.[10] 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 (see Appendix Figure A.1), there is no 'flash' or any warning sign that indicates the cameras' activity.

All recorded data are sent electronically to the local police, 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 manage the enforcement process.[11] The speed cameras

[9]On the benefits of ambiguity in an enforcement context, see also Lang and Wambach (2013).
[10]The number plate is retrieved by applying an automatic number plate recognition technology.
[11]The process is standardized, leaving no scope for discretion by police officers (Makowsky and Stratmann, 2009).

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. The tickets are sent to the car owners either via mail or electronically. As in many other countries, penalties are stepwise increasing in the speed (Traxler et al., 2018). Minor speeding offenses with a speed of 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 an intermediate speeding offense, with a speed 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 more than 40km/h above the limit, are handled according to a different procedure. Our analysis will not examine these offenses, which are very rare in our 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 rates do not increase either. Hence, the future 'price' of speeding does not increase with a speeding ticket; it remains constant. Second, tickets never result in driving licenses being revoked or suspended (or any jail sentences). Different from, e.g., Hansen (2015) or Gehrsitz (2017), punishment does not include incapacitation.

Third, the speed used to determine which penalty applies derives from a simple 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. A measured speed of, e.g., 73.85km/h is thus adjusted to 70km/h. Given this procedure, the cutoff for intermediate speeding is 23km/h in terms of *measured speed* above the limit. In the remainder of the paper, we will work with the precise speed measure, i.e., before the adjustment procedure is applied.

Fourth, when the speed cameras were set up, the local police decided to only send out speeding tickets 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]

Data. The city of Ricany provided us with data on the full universe of all 26 million rides recorded by the speed cameras from August 2014 through August 2018. For each car, we observe the exact time of entering and exiting a camera's zone, the measured speed (to the precision

[12] During 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.

[13] About 80% 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 and 2500–5000 CZK for minor and intermediate speeding case, respectively. 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).

[14] We 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 replies by the local authority are accessible online here.

of 1/1000 km/h), and an identifier for the specific zone. The data also include an anonymized (number plate-based) ID identifying a 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 – speeding and non-speeding rides – of each single car that was ever recorded by one of the five cameras. This clearly distinguishes our set-up from related empirical work.

The speed camera data were merged with administrative data on the enforcement of the speeding tickets. We observe the sending day for each ticket, whether it was sent by mail or email, and when it was received.[15] The data also include 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, decomposed for cars that never received a ticket and cars that did. The data set covers 26 million rides from over 1.3 million cars. Only a small fraction of cars received any tickets; specifically, 48,422 cars received over 56,000 tickets. The ticketed cars drive more frequently (84 rides on average compared to 16 for non-ticketed cars).

Our later analysis mainly uses two outcome variables: the *measured speed* and a *speeding* dummy which equals one if 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 the never-ticketed cars and 0.189 among ticketed cars. Note that the definition of the former variable includes all rides that violate the traffic law, irrespective of whether it is actually ticketed (i.e., speed exceeding 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 5.17 km/h below the limit) than the never-ticketed cars (6.00 km/h below the limit). As a third outcome variable, we use a (re)offense indicator which captures if 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 definition, the offense rate is higher (1.5%) among cars that ever received a ticket.

It is worth stressing that our data allow us to track cars but not individual drivers. Tickets are mailed to the owner of a car (who can differ from the driver) and we cannot distinguish different drivers that may share a given car (e.g., family members or different employees of a company). While our analysis at the level of cars likely capture 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 cars.

[15] The ticket is recorded as received when the addressee signs a delivery receipt (conventional mail) or opens the electronic mail with the ticket.

3 Theoretical Framework

This section illustrates the potential impact of speeding tickets in a simple framework that models optimal speed choices of a risk neutral driver. Note that the analysis can be easily extended to other domains where (i) the expected punishment is convex in the magnitude of a 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 not yet punished (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 $v(s_t, c_t)$.[16] $v(.)$, 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 $q^t(.)$ 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: $q^t(s) = \frac{\partial q^t}{\partial s} = 0 \, \forall s < \hat{s}, t$. The driver's problem is

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

and the optimal speed s_t^*, for a given expectation $q^t(.)$, is characterized by the first-order condition

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

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 $\mathrm{T}^t(s_\tau) \in \{0, 1\}$. Based on the past driving and ticketing experience, the driver might then update expectations:

$$q^t(s) = P\left(\{s_{t-1}, \mathrm{T}^t(s_{t-1})\}, \{s_{t-2}, \mathrm{T}^t(s_{t-2})\}, \ldots, \{s_0, \mathrm{T}^t(s_0)\}, \; q^{t-1}(s)\right), \tag{3}$$

where $P(.)$ 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

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

[16] We 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 $\vec{\mathrm{T}}(t, s_\tau) := (\mathrm{T}^t(s_\tau), \mathrm{T}^{t-1}(s_\tau), \ldots, \mathrm{T}^{\tau+1}(s_\tau))$ indicates if 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 expectations do not evolve over time. In this case, receiving a speeding ticket should not have any impact on subsequent speed choices. This is in contrast to predictions obtained for an imperfectly informed driver. For such a driver, a ticket provides new information regarding the probability of detecting a speeding offense (at a particular location and at a given speed s_τ) and the resulting consequences. We discuss alternative updating rules in the spirit of a reinforcement learning process, which yield different shifts in the expected costs $q^t(s)$ that imply different responses to receiving a ticket.

Consider first a driver that, after a ride with $\{s^*_{t-1}, \mathrm{T}^t(s^*_{t-1}) = 1\}$, 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 step right below s^*_{t-1}, the car would drive strictly below s^*_{t-1} even for 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 already very steep 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 thus drive very 'cautious', at most engaging in minor violations of the speed limit, even for 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.

Independently of whether a driver responds to a speeding ticket by updating expectations in a more fine grained or in a more coarse manner, the updating 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 nuances of updating have important implications, among others, for the ambiguity regarding the specific enforcement cutoff.[17]

[17] While the case for negative reinforcement (i.e., responses to speeding tickets) is fairly clear, one could also consider the possibility of positive reinforcement. This might occur when drivers' respond to 'successful speeding' (rides with $s^*_t > \hat{s}$) that did *not* result in a ticket. One might argue that, judged against 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 naive, to the extent that speeding in period t might result in a delayed ticket that arrives in any future period $\tau > t$. If such a positive reinforcement nevertheless occurs, drivers might become more 'optimistic' (in terms of $q^t(s)$) which would, cet. par., work towards increasing optimal speed.

3.2 The 'True' Enforcement Process

So far, our discussion 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 probability of triggering a ticket is equal to zero for any speed below the enforcement cutoff. For any speed above this cutoff, however, the chance of getting a ticket essentially equals one. Agents might learn about this enforcement cutoff. In fact, if drivers apply fine grained updating rules (as the one 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 would 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 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) ambiguous of the enforcement threshold contributes to larger behavioral responses to speeding tickets.

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).[18] This might be even more relevant as the second cutoff which is – in contrast to the enforcement cutoff – stipulated by the law and thus public. If the second cutoff would be unknown, however, the discussion from above regarding the scope for learning applies accordingly: whenever drivers update 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. As a benchmark case, one could consider the possibility that updating is insensitive to the level of the fines. However, a reinforcement logic would suggest that a stronger payoff loss (a higher fine) induces a stronger updating response.

With more fine grained updating in response to a low-fine speeding ticket, a larger penalty might indeed produce a more pronounced change in expectations. A case along these lines

[18] This 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.

is illustrated in the left panel of Figure 2. Under a more coarse updating, however, drivers might already respond to a low-fine ticket with a very conservative updating (see right panels of Figure 1). It is then questionable that a higher fine would induce any stronger behavioral responses. With more coarse updating, intensive margin variation in the experienced punishment might therefore not amplify the impact from speeding tickets on speed choices.

One refinement of the latter prediction builds on the possibility that the $q(s)$ function becomes 'more convex' after experiencing a higher rather than a low fine. Such a case is illustrated in the right panel of Figure 2. From the graphical illustration and the discussion of optimal speed choices above it follows that behavioural differences (in terms of different speeding levels) would only realize for sufficiently good speeding conditions c_t. We will test this prediction below.

The main behavioral predictions discussed above are summarized in Table 2. There are several additional dimensions worth discussing. Note first that, within our framework, any updating should produce an immediate effect. As long as there is no unlearning or forgetting, the response in speed choices after receiving a ticket should also be persistent. When drivers are constrained in their attention or cognition (Gabaix, 2019) – if they 'forget' or simply not have the risk of speeding tickets on top of their mind – there might be scope for 'unlearning': speeding tickets could serve as reminders that make the enforcement system salient. Receiving a ticket should trigger a temporary decline in speeding but the effect should decay over time. We might thus observe some 'backsliding' in the mid-run. Section 5 examines this possibility in an event study design.

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 (i.e., receiving or not receiving a speeding ticket); the cutoff that separates minor from intermediate speeding offense provides variation in punishment 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, among others, 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 and the date the maximum speed was recorded. If S_i is more than 14km/h above the speed limit ($D_i^1 = 1$), the car typically receives a speeding ticket. If S_i is more than 23km/h above the limit ($D_i^2 = 1$), the fine more than doubles:

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

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 sending date of a ticket that was triggered on that day. 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 the 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 if ride t of car i is 'treated': around the first cutoff ($k = 1$), treatment refers to a ride that happened 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 as well as cars with some foreign number plates will have $T_i^k = 0$).[21] On the other hand, this also reflects variation in the tickets' sending days: some tickets might be mailed days or weeks after the first ticket (for speeding offenses from that day) was sent. During the early phase of the outcome period, the rides of many cars will be untreated ($T_{it}^k = 0$), resulting in $T_i^k < 1$.

[19] Our 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.

[20] As 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).

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

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:

$$\begin{aligned} T_i^k &= \delta^k D_i^k + \kappa^k(S_i) + u_i, & (6) \\ Y_i &= \tau^k D_i^k + \lambda^k(S_i) + v_i, & (7) \end{aligned}$$

for both cutoffs, $k = \{1, 2\}$, and D_i^k as defined in (5). $\kappa^k(.)$ and $\lambda^k(.)$ 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.

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 the two coefficients one obtains the familiar 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 indicated 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, irrespectively 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, implicitly puts more weight on cars with more observed rides (i.e., more frequent drivers). This will yield effects for an *average ride*.

Our main analysis of the first cutoff will include 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., $0 < S_i < 23$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 $14 < S_i < 43$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. Let us first provide graphical evidence on the treatment discontinuity around the first cutoff. Figure 3 plots local linear fits, confidence intervals and binned averages for the treatment rate T_i^1, the share of 'ticketed rides' (after having received a ticket) during the outcome period. The figure – which covers the range of assignment speeds 10km/h below and up

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

to 9km/h above the enforcement cutoff[23] – shows a clear discontinuity in the extensive margin of punishment: at the cutoff, the share of treated rides jumps from marginally above zero to roughly 80%. The underlying estimates (documented in Column (1) of Table A.3) indicate a 78.7 percentage-point (pp) discontinuity in the share of treated rides.[24] As discussed above, the share of treated rides for S_i above the cutoff is below one: emergency cars (and some cars with foreign number plates) are exempted from the enforcement process and some rides during the outcome period may 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 authority does not systematically prioritize sending out tickets earlier for offenses with a higher speed (within the range of minor offenses). We will return to this observation below.

Sorting. To validate if the treatment discontinuity offers as-good-as-random local variation, there must be no sorting of cars below the cutoff.[26] There are several institutional features which make sorting appear implausible in our context. First, the actual enforcement threshold is not publicly known. As pointed out above, the cutoff is not prescribed by the law but was determined by the police once the speed cameras started working (see also 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.[27] Optimal targeting is further complicated by the applied tolerance rule (see 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 discussed in Section 3, however, there might be nevertheless 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 on sorting below the cutoff (see Figure A.2).[28] The data do not indicate any bunching, not even 'imprecise' one. Moreover, there is absolutely no evidence on the emergence of heaping over time. This last statement holds for the sample of all cars (Figure A.3) as well as for 'regional cars'

[23] Recall that the second cutoff (23km/h above the limit) is 9km/h above the first one (14km/h above limit).

[24] Throughout the paper we report bias-corrected RD estimates with robust variance estimators (Calonico et al., 2014), implemented with the 2018/09 version of 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.

[25] Figure 3 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 $S_i <$ 14km/h is speeding with more than 14km/h above the limit during the outcome (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.

[26] As pointed out by Lee and Lemieux (2010), the RDD would be still valid if drivers can only imprecisely choose the running variable, S_i.

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

[28] Panel (a) of Figure A.2 provides weak evidence on a minor increase in the density on 'the wrong side' of the cutoff. The estimate, however, is economically negligible and sensitive to bandwidth choice.

(number plate from the local regioan and above-median driving frequency; Figure A.4). Note that these observations are consistent with the implications from coarse updating (see Section 2 and Table 2).

Balance. Next, we examine if 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 plate 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 this occasion).

4.2.2 High-fine cutoff

Treatment. Evidence on the treatment discontinuity around the high-fine cutoff is provided in Figure 4.[29] The share of 'high-fine treated' rides – rides after having received a ticket with a high-fine, T_i^2 – discontinuously increases by 81.4 percentage points (Column (1) in Table A.4). Hence, at the second cutoff there is a strong discontinuity in the exposure to high- rather than low-fine tickets.

Sorting. As for the enforcement cutoff, we do not detect any evidence on 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. Similar as above, we also tested if bunching would emerge over time. The data do not provide any evidence on this case, 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 scope for one complication that relates to the way speeding tickets are sent out. Figure 3 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 quicker. In turn, this might 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 (see 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.

[29] Recall 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 4.

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 roads covered by the speed cameras are commuting routes that are difficult to circumnavigate. It is thus not surprising that we find no evidence of cars either reducing their driving frequency or stopping to drive in response to speeding tickets (see Columns 1 and 2, Table A.5).[30]

Next, we turn to reduced form evidence on speeding responses. Figure 5 show pronounced discontinuities in the cars' speeding rates and the mean speed measured during the outcome period. Cars with an assignment speed marginally above the enforcement cutoff have a 8.1pp lower speeding rate and their mean speed is about 1.35km/h slower (see Columns (2) and (4), Table A.3).[31] Below we will see that this decline in the mean speed masks stronger responses at the top of the speed distribution.

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 a speed 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).[32] This decline is more pronounced at the top end of the speeding distribution: when we estimate the effect of the speeding ticket on a car's speed at the median, the 75th- or the 90th percentile of its' speed distribution, we observe an increase in the absolute (from 1.31 to 1.77km/h) and the relative effect size (from 2.8 to 3.4%; see Columns 4–6, Table 3).

The (reduced-form) effect from receiving a speeding ticket on the speed distribution is also depicted in Figure 6. The dashed, red line illustrates 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 3). 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 – one notices a clear shift in the distribution. Among cars that are marginally above the cutoff, rides with speed above the limit are observed less frequently. This missing mass is mostly shifted

[30] In 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 the 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).

[31] Recall 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.

[32] This estimate slightly differs to what would be obtained from simply deflating the reduced form coefficients from Table A.3 by the treatment discontinuity, as the MSE-optimal bandwidths for the Wald estimators (that jointly estimates δ^k and τ^k from (6) and (7)) are different to those for the reduced form estimates.

towards the mode of the distribution, which is (for both groups) roughly 3km/h below the speed limit.

The shift in the speed distribution is also illustrated in Figure 7. The latter figure plots the relative change in the mass between the two speed distributions from Figure 6, in 7km/h bins above and below the actual speed limit.[33] In line with the strong drop in (re-)offenses reported above, the figure illustrates an approximately 50% drop in the mass of rides with a speed of 14–21km/h above the limit. Consistently with a coarse updating response, however, one observes 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. Note that this observation is inconsistent with a very nuanced, fine grained updating of expectations.

To wrap up, the basic estimates as well as the graphical evidence documents 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 basic estimates we first consider alternative bandwidths. Figure A.8 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 up to 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 has indeed an impact on sample size and composition (with shorter periods, we tend to observe fewer infrequent drivers), our estimates are remarkably stable for different combinations of a and f values. This point is documented in Figure A.9, 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 high robustness w.r.t. these two values foreshadows two results from below. We will see, on the one hand, that type heterogeneity plays a modest role. On the other hand, 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 ($f = 6$) solely matters 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.[34] Compared

[33]To compute relative changes, we normalize the observed mass in the groups marginally below and above the cutoff. Due to the extent that (i) only 80% among the latter group are treated and that (ii) pre-treatment speed distribution (i.e., for rides observed during the assignment period) are, by construct, different between the two groups, the numbers from Figure 6 represent approximations for the treatment impact on the speed distribution.

[34]The estimates are again robust w.r.t. different a- and f-periods, see Figure A.12.

to the results from the collapsed analysis (see Table 3), we get slightly smaller point estimates, in particular for the effect on average speed. As we will further discuss 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, we shall note that 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, differential responses are not very pronounced (see Panel A of Tab. 5). (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) as compared to the other cars. These findings must be interpreted with caution, however, as the 'local' number plates include a relatively large area beyond Ricany.

In a further step, we also compare the effects on rides occurring under more or less favourable traffic condition (c_t), as captured by the traffic density (measured by the time gap to the next car ahead). Consistent with the comparative static implications from our formal 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 by 15pp (37%). Under bad conditions, these estimates are much smaller (5pp drop in speeding rate and 0.75km/h decline in mean speed; see Panel A in Table A.8). These findings have to be interpreted with caution as driving conditions in the outcome period are potentially shaped by the choice when to drive.[35]

4.3.2 Punishment at the intensive margin (high-fine cutoff)

Next we turn to the second cutoff, which provides variation in punishment at the intensive margin. Similar as above, we first examine whether receiving a high-fine (as 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 turn to our two main outcome variables. Figure 8 does not indicate any clear discontinuities, neither in the speeding rates nor in the mean speed (see the corresponding reduced form estimates in Columns 3 and 5,

[35]Using day-of-week and hour-of-day indicators for rides in the outcome period, however, we do not find any evidence on speeding tickets shaping the timing of rides.

Table A.4). The Wald estimates from Table 6 – which are based on a much smaller number of observations as compared to those at the first cutoff (16 rather than 225 thousand cars) – include no statistically significant estimates either. While the evidence suggests that the average car does not respond differently to a high- or a low-fine ticket, all estimates are negative and some effect sizes seem large (at least in relative terms). We thus explore the sensitivity of these insignificant findings.

Concerning the different bandwidth choices, the estimates turn out to be robust (see Figure A.11). When we focus on shorter outcome periods, however, we tend to find weakly significant effects on the mean speed (but not on the speeding rate; see Figures A.12). Next we replicate the estimates at the level of rides. Despite boosting the number of (clustered) observations, the ride-level estimates do 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 more, we find no statistically significantly differences in the responses to tickets with higher fines (see Table A.9).

The predictions from the analytical framework presented in Section 3 suggests that, under coarse updating, intensive margin variation in penalties does not necessarily amplify the impact from receiving a ticket. A notable exception emerges when higher penalties produce 'more convex' expectations (as illustrated by the $q(s)$ functions in the right panel of Figure 2). In the latter case, the intensive margin variation in punishment would only induce behavioural responses when driving conditions c_t are sufficiently good. This prediction is examined in Table 7, which presents (car-level) estimates for good and bad traffic conditions (as measured by the time difference to the next car in front at the entry of a speed camera zone).

For rides observed under relatively dense traffic, we estimate economically small and statistically insignificant effects (Table 7, Columns 4–6). Under 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 speed drops by another 1.5km/h (–3%) and 2.1km/h (–4%), respectively (see Panel A, Columns 1–3 of Table 7). As 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 facing 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, whereas other estimates become smaller and turn insignificant again (Panel B of Table 7). Hence, the findings only provide weak support for the more convex updating illustrated on the right hand side of Figure 2. As pointed out above, however, one has to cautiously interpret these results, as drivers might (conditional on the high-fine treatment) select into good or bad driving conditions. Overall, the evidence suggests that the variation in fines seems to play a minor role for the way drivers update and respond to tickets. The effect from intensive margin variation in fines seems to be 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 high-frequency nature of our data to examine how quickly drivers respond and how long-lasting the effects are. We can also compare the within-estimates from the event study, which yield an average treatment effect on the treated (ATT), with the LATE obtained from the RDD.

5.1 Design and Sample

For each car receiving 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. We further focus on low-fine tickets, triggered by speeding between 14 and 23km/h above the speed limit. This allows for a meaningful comparison of the event study ATT with the LATE at the enforcement cutoff (see Section 4.3.1).

Figure 9 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[36] which satisfy the sample conditions described above. The time axis is defined such that week zero is the last week before the 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.[37] The graphs indicate strong and persistent treatment responses: after receiving the 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.

However, Figure 9 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. The pattern simply reflects 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 three weeks prior to receiving a ticket. This explains the pronounced increase in speeding observed in the raw data. A naive estimation that would include these humps would then overestimate the impact from tickets. (A formal discussion of this point is provided in Appendix B.3.)

To deal with the issue, we exclude the trigger observations from our analysis. (Note that this is a fairly conservative approach, as the trigger observation is a relevant observation for a car's behavior prior to the ticket.) The effect from this exclusion is illustrated by the lines indicated with triangles in Figure 9: the massive humps disappear and pre-ticket trends are modest. The final analysis-ready sample then includes 626,430 rides from 16,407 cars for their first (low-fine) ticket event. We analogously define a sample for the first ticket event punished with a high-fine. Later we will also examine cars around their second ticket event.

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

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

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

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

where Y_{izt} is a speeding outcome of car i observed at speed camera zone z and 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 schooldays/holidays (λ_{dz}) as well as hour of the day dummies (λ_{hz}). As driving patterns differ between zones, all these time-specific dummies are interacted with the zone dummies. We also include a vector of variables capturing the driving conditions for a given ride (X_{izt}). The vector includes, among others, a set of dummies that non-parametrically capture the traffic density at ride ict as well as weather variables (temperature, precipitation, sunshine intensity, measured at a 10-minute frequency).[38]

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).[39]

5.2 Event Study Results

5.2.1 Response to Punishment

Panel (a) of Figure 10 plots the estimated coefficients and confidence intervals for the binary outcome speeding. 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 degree, in the 3rd) week after facing the ticket. The effects are very 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 slight but statistically insignificant downward trend). Compared to the pre-ticket baseline, the 10pp drop implies a 37% reduction in the speeding probability.

Panel (b) of Figure 10 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

[38]To capture the strong influence of the traffic situation, X_{izt} includes dummies for whether the car 'ahead' of car i (at the ride at time t and 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 included the total number of cars passing the 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 weather conditions are practically identical at all five speed camera locations.

[39]Clustering only at the level of cars yields similar results.

additional reduction in later (mainly the 2nd) weeks. Over the 20 weeks, the estimated effect size varies but remains approximately flat (again, with a slight downward tendency) in the range of 1.2–1.4km/h below the baseline speed level. Relative to the baseline average of 44.86km/h in the last pre-ticket week (see bottom panel of Table A.11, Column 2), the 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 the mid-run. In terms of absolute and relative effect sizes, it is worth noting that the ATTs obtained from the event study design are very similar to the LATEs found in the RDD analysis from above. Note further that we do not find any evidence on 'backsliding': speeding outcomes do not revert towards the pre-ticket levels. This clearly rejects the idea of 'unlearning'. The results are consistent with the learning channel and the large drop in speeding rates supports the notion of coarse updating examined in Section 3. Drivers seem to learn about the consequences of speeding, update their expectations such that $q^t(s)$ increases steeply just past the speed limit, and adjust their speed choices accordingly.

Further evidence that supports coarse updating is provided by Figure 11, which depicts the effect on the speed distribution. It is analogous to Figure 6 from above, except that it is based on a within-car comparison: the two lines compare the same 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 shifted predominantly to the speeds about 5km/h below the speed limit (which is also the mode of the pre-ticket distribution). Such a shift is inconsistent with fine-grained updating which would result in an increase in the mass below the enforcement cutoff.

Figure 10 indicate that effects are persistent over a 20 weeks period. To explore whether there is no backsliding in the long-run, we estimate an alternative specification which (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 sizable number of rides (991,333).

The estimated coefficients on the monthly dummies are plotted in Figure 12 (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 not evidence on backsliding. On the contrary, the effect size slightly increases over the two years outcome window. (This seems to reflect the general decline in speeding observed at 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. 17). 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, at least for the average driver, with no 'unlearning' or declining salience of the enforcement regime over time.[40]

5.2.2 Response to Higher Fines

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 small-fine tickets – are presented in Figure 13.

For the speeding rate, the average effect sizes are virtually identical between cars receiving a high- or a low-fine ticket. 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 13). 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 44.86 km/h) in the high-fine sample.[41] Similar to the RDD analysis, these estimates do not provide much evidence of an additional effect from an intensive margin increase in penalties – at least on average. (In future work, we plan to explore differential patterns for rides under different driving conditions.)

5.2.3 Heterogeneity Analyses and Extensions

Analyzing if and how the effects vary across different types of car owners, we follow the RDD analysis and first compare frequent and infrequent cars (as measured by the cars' pre-ticket driving frequencies). The results, which are presented in Figure A.13, 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.14 compares cars according to the number plate region of a car. Consistent with the RDD results reported in Table 5, we observe smaller effects for cars from Prague and the local region relative to cars from other regions of the country. The differences, however, are in general statistically insignificant.

The event analysis, which in contrast to the RDD focusses on cars that receive a ticket, enables us to explore further dimensions of heterogeneity. Among others, we observe whether a ticket was sent to a physical person or a 'corporation'.[42] If the car owner is a private person, the

[40] This might be due to the visibility (and stable functioning) of the speed cameras, which serve as constant reminders.

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

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

individual (who is also in charge of paying the fine) will directly learn about the speeding ticket. For cars owned by corporations, in contrast, there might be more frictions in the learning process. Corporations might pay the tickets on behalf of their drivers without informing them. In case of multiple drivers sharing one car, it might be hard to identify the responsible driver. Even if the message reaches the (relevant) driver, it is unclear whether all 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 a speeding ticket than drivers of cars owned by corporation. Figure 14 provides some support to 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 of these differences weakly 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 should be particularly relevant for larger corporations).

A final dimension of heterogeneity is based on whether the car owner did or did not pay the speeding ticket (within 90 days).[43] 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 15. The estimates indicate that cars who pay the ticket are slowing down much more strongly. Those who do not pay the ticket nevertheless adjust their driving behavior. During the first three weeks, the speeding rate among the former group drop by 8–11pp. Among the latter, the drop amounts to a mere 3–5pp. Over time, this gap narrows but does not fully disappear in later weeks. A similar (but less precisely estimated) pattern is also observed for the level of speed.

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 during highly congested traffic conditions, and by 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.

A 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 (and potentially vanish after many repetitions). To account for this, we replicate the event study with treatment dummies defined by the order of rides: we sort rides within each car and then define dummies grouped over intervals of five rides before and after the 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 from this exercise are reported in Figure 16, where the omitted category is now given by the last five rides before the ticket.

[43]The average time from receiving the ticket to payment is 10 days, conditional on the payment being made. 15% of car owners do not pay their tickets by that time.

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 before receiving a ticket. In addition, we also observe a clear increase in the effect size between the 5th and the 20th post-ticket ride. Both observations, however, are partially shaped by changes in sample composition, as observations further away from the ticket are increasingly composed of cars with higher driving frequency.

Second Ticket. So far, we have focused on the impact of the first speeding ticket. However, some cars receive a second ticket later on. Among the 33,016 cars that received at least one ticket, 17.52% get a further ticket.The probability of receiving a second ticket is significantly higher for cars owned by corporations. It is also strongly positively correlated with the measured speed of the first ticket, suggesting that cars with a stronger taste for speeding select into this small group of roughly 6 thousand cars (out of a total of 1.3 million, see Table 1).

Independently of the selection process, we can estimate behavioral responses around the second ticket event. We focus on 2,566 cars that were punished for their first speeding offense, later received a second (low-fine) ticket and have at least one non-trigger observation during the pre- and post- windows around the second ticket event. The results are presented in Figure 18. While the sample is highly selected and much smaller, the baseline pre-ticket mean outcomes are similar to our main sample. The responses to the second ticket are qualitatively similar to the responses to the first ticket: on average, we observe again an immediate, large and sustained reduction in the speeding probability and the speed. Quantitatively, however, the effects are about 20% smaller than our basic estimates for the first ticket. One interpretation of these findings is that these cars represent fine-grained updaters: drivers shift their expected costs $q^t(s)$ only slightly, which in turns leads to a modest decline in speed. Eventually, they get a second ticket which is followed by another small adjustment in q^t and optimal speed (relative to the average response to the first ticket). In follow-up work, we plan to explore this interpretation in more detail, by examining the responses of this subgroup of cars to the first ticket.

6 Further Results

6.1 Narrow or Broad Learning?

The results from the RDD and the event study coherently document drivers' responses to facing punishment. The evidence rejects the case of no-updating and supports the learning and (coarse) updating framework from Section 3. This subsection now explores whether drivers' learning is more 'broad' or 'narrow'. More specifically, we ask whether drivers solely update $q^t_z(s)$ – the expected costs of speeding in the camera zone z that triggered the ticket – or whether they would update expectations $q^t_\ell(s)$ 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* zone or in *other* zones. Figure 17 plots event study estimates for this comparison.[44] We observe clear behavioral responses in the same zone but also at the 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 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 along 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). Results are similar when we constrain the sample to cars that are observed in both, the same and at other zones (Columns 3–4). For the relative declines in 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 estimates from the event study design are similar (see Table A.16).

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

To tackle this point, we exploit that our data contain the exact time when a car 'exits' from (the endpoint of) one speed camera zone and 'enters' into (the starting point of) another one further down the road. Based on this time gap, we can learn about cars' speed in the unmonitored stretch in between: the faster a car drives, the quicker it will enter into the second camera zone. Theoretically, one can derive three very different predictions about the impact of a speeding ticket: a first hypothesis is that, while cars are updating $q_z^t(s)$ at different zones z (see above), they would not do so at roads not covered by speed cameras. In this case, we should not see any change in the time spent on the unmonitored part of the road. A second hypothesis is that the learning spills over and indeed results in a broad adjustment in expected fines, beyond the zones covered by the cameras. One would thus expect a drop in the inter-zone trave time. (This could also happen because 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., loosing time) in the speed camera zone, they might want to 'catch up' by driving faster on the unmonitored part of the road.[45] We should then observe faster inter-zone rides in response to a speeding ticket.

[44] Note that the camera that triggered the first ticked varies between cars.

[45] In terms of the model from Section 3, the marginal benefit of speeding, $\partial v(.)/\partial s_t$, might increase after having slowed down before.

Among the five speed camera zones there is only one combination where, after exiting one zone, one could enter (after a left turn) another camera zone. This 'unmonitored trip', however, is not very frequently observed. Moreover, drivers encounter a traffic light on the way, which introduces sizable variation in the travel time on the 1,080 meters between the exit of the first and the entry into the next zone. One can nevertheless use the measured time to compute the average speed on the trip. We then apply the RDD strategy from above (focussing on the variation around the enforcement cutoff) and estimate effects on inter-zone rides.

The estimates, which are presented in Table A.10, reject the third, 'catch-up' hypothesis from above. Columns (1) and (2) present car-level estimates for the (reduced form) effect on average and on the top- (90th percentile) speed. For the mean speed, we obtain a negative but imprecisely estimated effect.[46] For the 90th percentile speed, however, we get a relatively large and weakly significant negative coefficient. Turning to ride-level estimates, Column (3) indicates a smaller, insignificant effect on speed.[47] Using the log travel time as a dependent variable, however, the effect is again significant and the positive sign indicates that cars slow down (take more time). In future work, we might be able to use a larger set of data stemming from new speed cameras. This should allow us to replicate this type of analysis and to increase power.

To wrap up, our analyses provide evidence that is consistent with the first and weakly supportive to the second hypothesis from above: some cars seem to reduce their speed on the unmonitored part of the road, too. Together with the results from the between-zone comparison from above, the evidence supports the notion of a broad learning (and behavioral) response to law enforcement.

6.2 Treatment Spillovers

In a last step, we analyze 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, under dense traffic conditions, such spillovers might reach beyond the next car in a line. In addition, we also explore spillovers on the car ahead. Given that the ticket makes an (otherwise aggressive) car drive slower, being less pushy might also affect the car in front of the 'ticketed' car.

To evaluate such spillovers we identify different *car groups* g in our data. In particular, we consider lines of two or more cars which all enter a camera zone within 5 seconds to each other. The first car (the 'start' of a line) is required to enter the measurement zone at least 10 seconds after the previous car. The last car of a line (which marks the 'end' of a group) enters the zone more than 5 second ahead of the next one. Based on these definitions one can then zoom into the sequences with different lines.[49] We study the responses of cars (rides) in position j in a given

[46] Presumably due to the traffic light, the average speed for this trip is far below the speed limit.

[47] See Section 4 for a discussion of the differences between (unweighted) car- and (weighted) ride-level estimates.

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

[49] The 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

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 the compare the *direct* treatment effects for different positions j.

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

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

The key parameter, τ_j^ℓ, measures the (reduced form) impact of a car in position ℓ that has an assignment speed above the enforcement threshold on rides of 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,

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

where the treatment dummy $T_{\ell gt}$ indicates if ride t of the car in position ℓ in group g is 'ticketed' (i.e., occurred after receiving a speeding ticket). Two remarks are at place 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 $j = \ell$, the estimates are conceptually analogous to the ride-level estimates presented above (see Table 4)). For $j \neq \ell$, the RDD exploits variation in *other* cars' assignment speed $S_{\ell g}$ rather than the 'own' S_i. Second, we estimate β_j^ℓ separately for all rides observed in position $j \in \{1, \ldots, 5\}$. 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.

Results from 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 or five and more cars in a line. The different columns present effects on the j = 1st, 2nd, ..., 5th 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 a direct treatment effect 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 above (see Column (1) and (3) in Table 4). In addition, the estimates also reveal meaningful spillovers on the 2nd and 3rd cars within these groups. In lines of three cars (Panel b), for instance, the speeding rate among the first (i.e., the 'ticketed') car drops by 7.4pp, and by 5.5pp and 5.1pp for the 2nd and 3rd car, respectively. A similar pattern is observed for the driving speed.

redundant and only serves to illustrate (consistently with the notation from Section 4) that we focus on ride-level rather than collapsed outcomes.

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

[51] The 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 position ℓ and for different group sizes (number of cars in a line).

Next, we consider groups of cars where the one in position $\ell = 2$ is ticketed (Columns 6 – 10). In addition to direct effects on cars in position $j = 2$ we again find some evidence on 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 scope for speeding of the treated car is already constrained by the mere fact that this car is observed in position $\ell = 2$ of a line moving in relatively dense traffic.[52] Despite that, we obtain some weak evidence on 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. Similarly 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, the estimates provide evidence on treatment spillovers. Especially for cases where the leading car of a line is ticketed, the speeding ticket also induces a decline in speed and speeding rate among (at least) the next two cars in the line. In addition to these (partially mechanical) backward spillovers, however, we also find some evidence on 'forward spillovers' on the car ahead. The latter spillover suggests that ticketed cars would have otherwise 'pushed' the one ahead to drive faster.[53]

The implications from 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, if these spillovers are all positive from a social welfare perspective (e.g., associated with lower noise and CO_2 emissions): 'backward spillovers' are responses to cars who slow down in front, which could – in case of abruptly braking cars – in principle increase accident risks.

7 Concluding Summary

Based on unique data that cover driving histories of 1.2 million cars over several years, we 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 of the speeding rate by a third and a 70% drop in re-offending. A doubling in the speeding tickets' fines has only limited additional effects. Event study estimates, which

[52] For, 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 leading cars within a line – which, by our definition of car groups, have a free road ahead – drive faster than the following cars. For instance, within lines of four cars (with the leading one being ticketed), the speeding rate of (marginally untreated) cars in position $j = 1, \ldots, 4$ 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.

[53] Sensitivity 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 at entering the camera zone is between 3–8 (rather than < 5) seconds).

confirm all LATEs from the RDD, further show that the responses are immediate and persistent over several years. Adjustments in speeding are observable in different speed camera zones and there is no evidence on compensatory speeding on un-monitored parts of the road. Instead, the data indicate spillover effects on untreated cars, which reduce their speed, too. Given that the WHO (2018) considers effective speed management policies as the 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 to isolate such learning-induced deterrence effects. The data are consistent with a coarse, discontinuous updating of priors (and reject a fine-grained way of updating). The persistency of the effect further rejects the interpretation of the findings in terms of temporary salience responses of agents with limited attention.

The 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 to numerous domains where offenders face ambiguity about the red line at which an authority's 'tolerance' of illegal behavior ends and where enforcement starts (e.g., petty theft, minor drug possession, public nuisances, tax evasion). It is up to future research to examine to which domains this applies.

References

Advani, A., W. Elming, and J. Shaw (2019). The Dynamic Effects of Tax Audits. CAGE Working Paper No. 414.

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.

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.

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. Giaccobasso, 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 and Timely Compliance: Experimental Evidence from the Enforcement of Speeding Tickets. Hertie School of Governance, Mimeo.

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.

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.

Kolm, S.-C. (1973). A note on optimum tax evasion. *Journal of Public Economics 2*(3), 265–270.

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. and A. Wambach (2013). The fog of fraud — Mitigating fraud by Strategic Ambiguity. *Games and Economic Behavior 81*(C), 255–275.

Lee, D. S. and T. Lemieux (2010). Regression Discontinuity Designs in Economics. *Journal of Economic Literature 48*(2), 281–355.

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. Terlizzese (2019). Leave the Door Open? Prison Conditions and Recidivism. Collegio Carlo Albert, Mimeo.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics 142*(2), 698–714.

Mobius, M. and T. Rosenblat (2014). Social Learning in Economics. *Annual Review of Economics 6*(1), 827–847.

Mueller-Smith, M. (2015). The criminal and labor market impacts of incarceration. University of Michigan, Mimeo.

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, Mimeo.

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.

Tang, C. K. (2017). Do speed cameras save lives? LSE Research Documents on Economics No. 86567.

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.

Figures

Figure 1: Updating in response to a speeding ticket

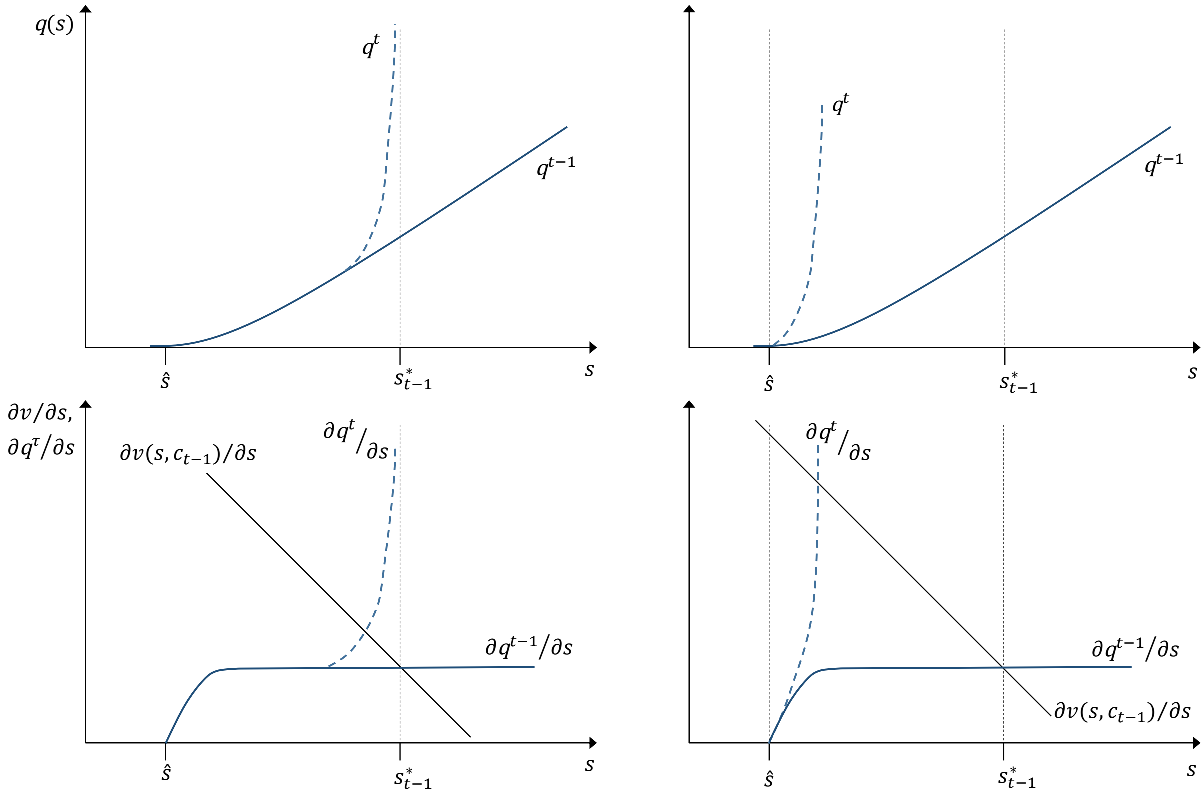

Notes: The figures illustrate possible updating responses to receiving a speeding ticket in period t after a ride at speed s_{t-1}. The figures on the left illustrates the case of a fine grained way of updating. The figures on the right hand side consider a more coarse updating response. The two panel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

Figure 2: Updating in response to low- vs high-fine speeding ticket

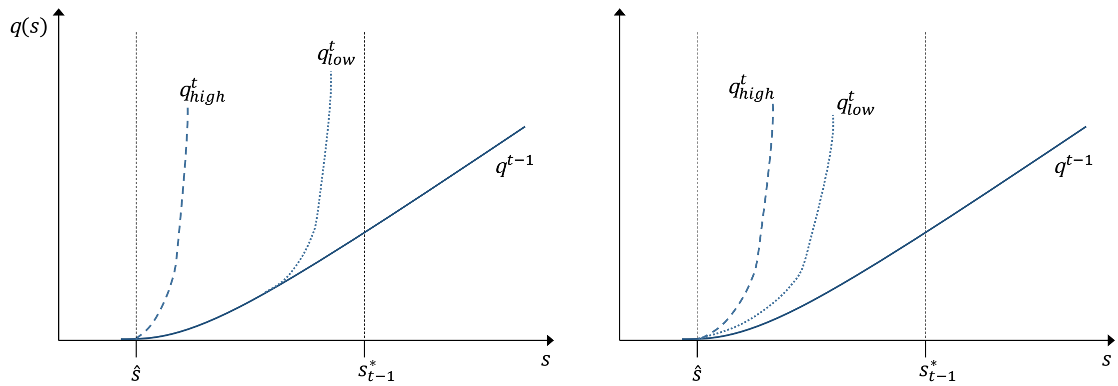

Notes: The figures illustrate possible updating responses to receiving a speeding ticket with either a low (q^t_{low}) or a high fine (q^t_{high}). The left figure captured the case where a higher fine results in a larger updating of expectations. The right figure depicts the possibility that, for a similarly coarse updating, the higher fine increases the convexity of q^t_{high} (relative to q^t_{low}).

Figure 3: Share of 'treated' rides, enforcement cutoff

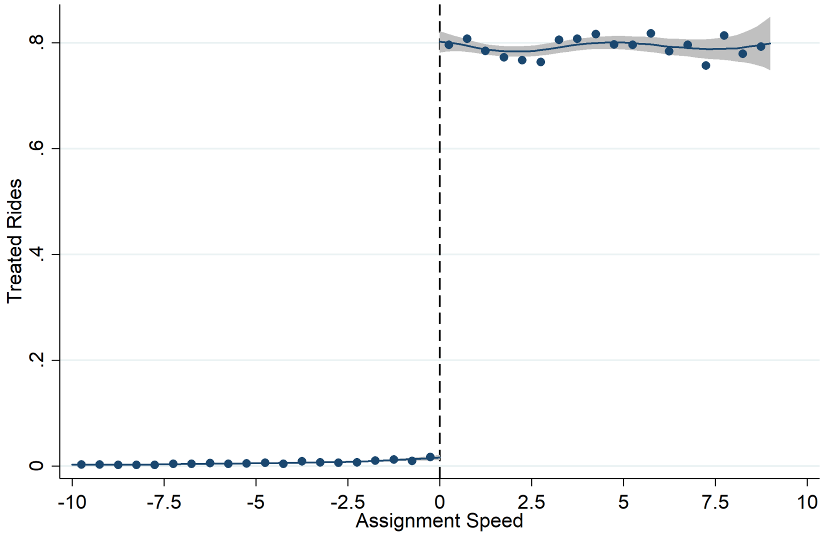

Notes: The figure presents the cars' share of *ticketed rides* T_i^1, i.e., rides after receiving a speeding ticket (relative to all rides in the outcome period), around the enforcement cutoff (1st cutoff). The assignment speed, S_i, is normalized relative to the cutoff (14km/h above the limit). Local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

Figure 4: Share of 'high-fine treated' rides, high-fine cutoff

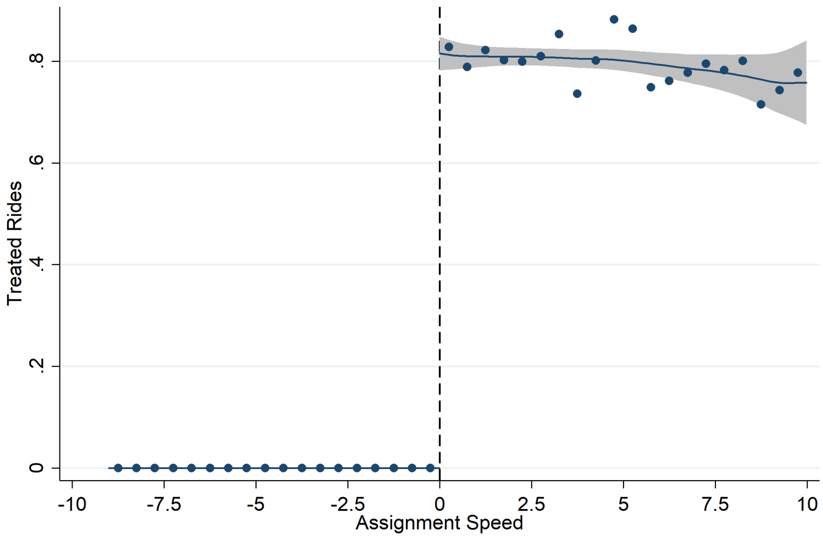

Notes: The figure presents the cars' share of *high-fine treated rides* T_i^2, i.e., rides after receiving a high-fine speeding ticket, around the high-fine cutoff (2nd cutoff). The assignment speed, S_i, is normalized relative to the cutoff (23km/h above the limit). Local linear estimates (with a MSE-optimal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

Figure 5: Discontinuities in outcomes at the enforcement cutoff

(a) Outcome: Speeding

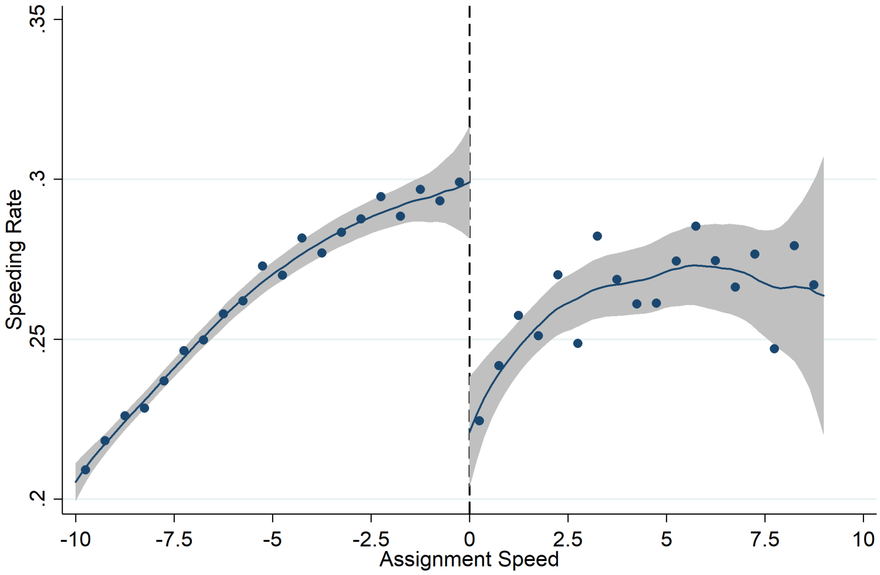

(b) Outcome: Speed

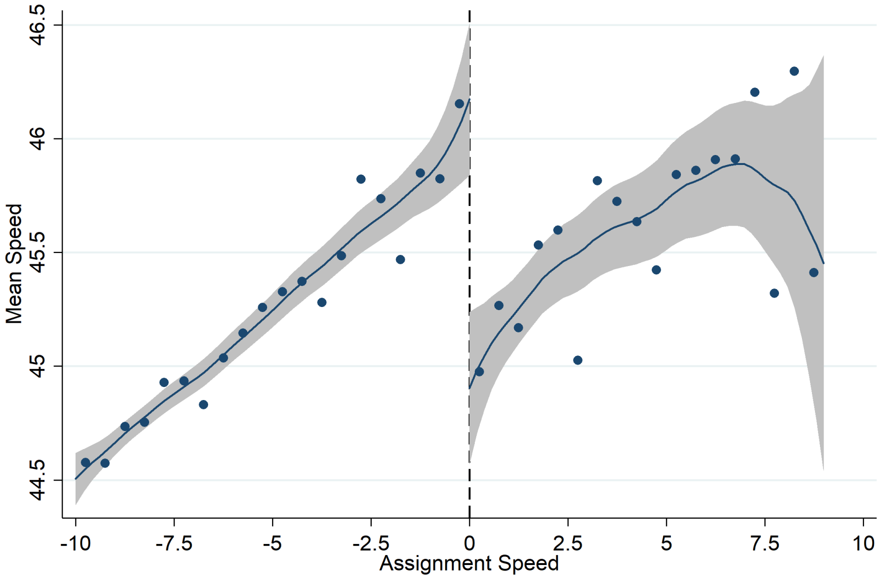

Notes: The figures present speeding rates (Panel a) and the cars' mean speed (b), i.e., the car's average share of rides above the speed limit, around the enforcement cutoff. The assignment speed, S_i, is normalized relative to the cutoff (14km/h above the limit). Local linear estimates (with a MSE-optimal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

Figure 6: Change in the speed distribution (enforcement cutoff)

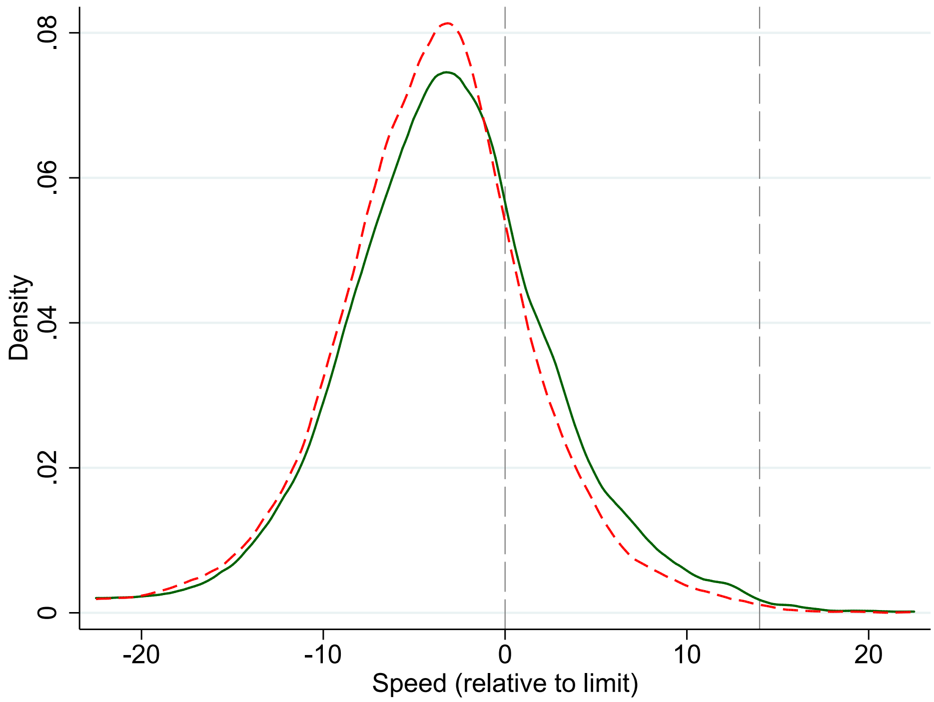

Notes: The figure illustrates speed distributions during the outcome period (with the measured speed normalized relative to the speed limit). The solid, green line captures the distribution for all rides from cars with an assignment speed S_i within a 0.5km/h range *below* the enforcement cutoff (i.e., 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 7: Relative change in speed distribution (enforcement cutoff)

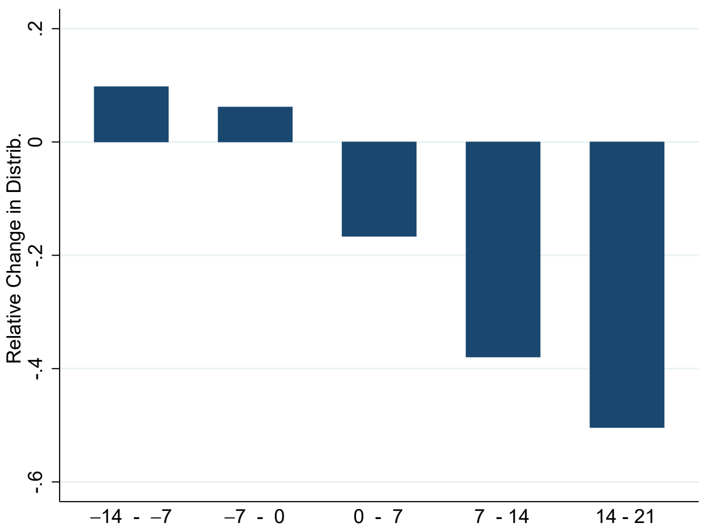

Notes: Based on the distributions displayed in Figure 6, we computed the difference in the speed distribution among cars with an assignment speed S_i within a 0.5km/h range *above* the enforcement cutoff relative to the distribution among cars with a marginally lower assignment speed (within a 0.5km/h bin *below* the enforcement cutoff).

Figure 8: Discontinuities in outcomes at the high-fine cutoff

(a) Outcome: Speeding

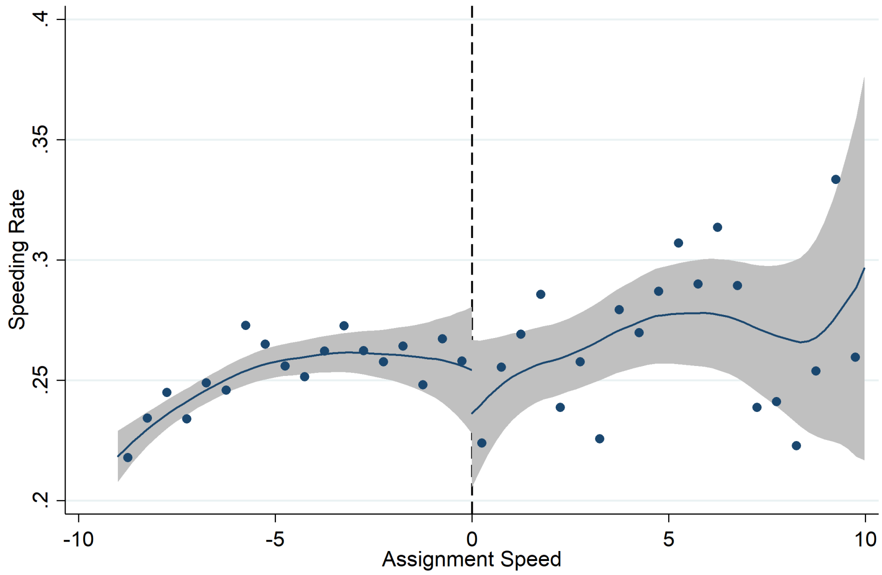

(b) Outcome: Speed

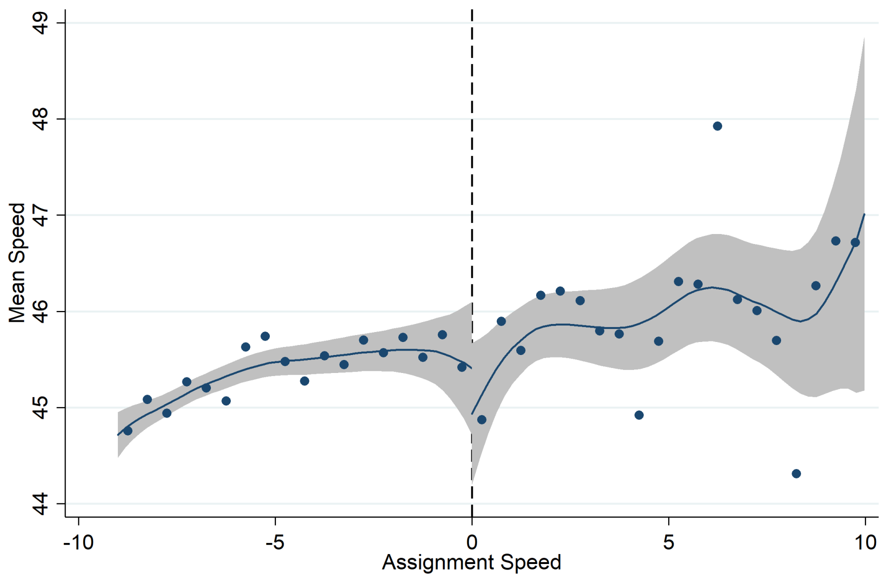

Notes: The figures present speeding rates (Panel a) and the cars' mean speed (b), i.e., the car's average share of rides above the speed limit, around the high-fine cutoff. The assignment speed, S_i, is normalized relative to the cutoff (23km/h above the limit). Local linear estimates (with a MSE-optimal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

Figure 9: Event study: plot of raw data

(a) Outcome: Speeding

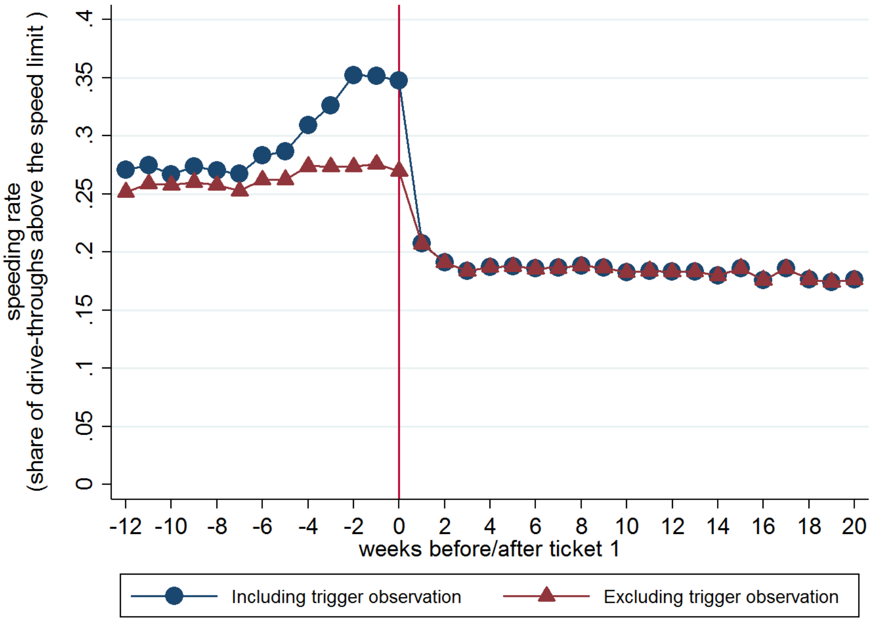

(b) Outcome: Speed

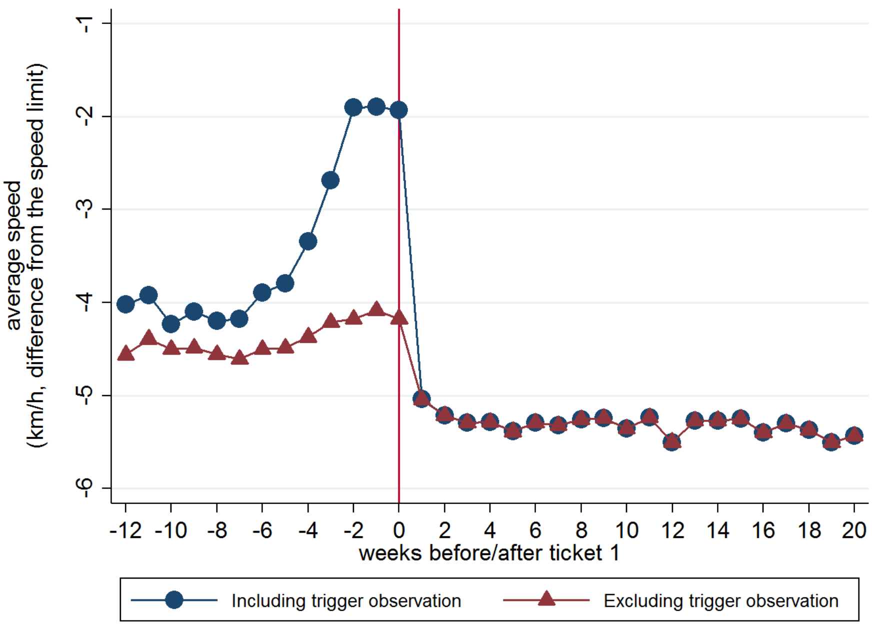

Notes: The figure plots speeding rates (Panel a) and the observed mean speed (b) in weekly intervals before and after receiving the first ticket. The sample includes cars around the first ticket event that faced a low fine. Cars included in the sample have at least one observation during the pre-ticket period (other than the trigger observation) and at least 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 10: Event study estimates: responses to the low-fine ticket

(a) Outcome: Speeding

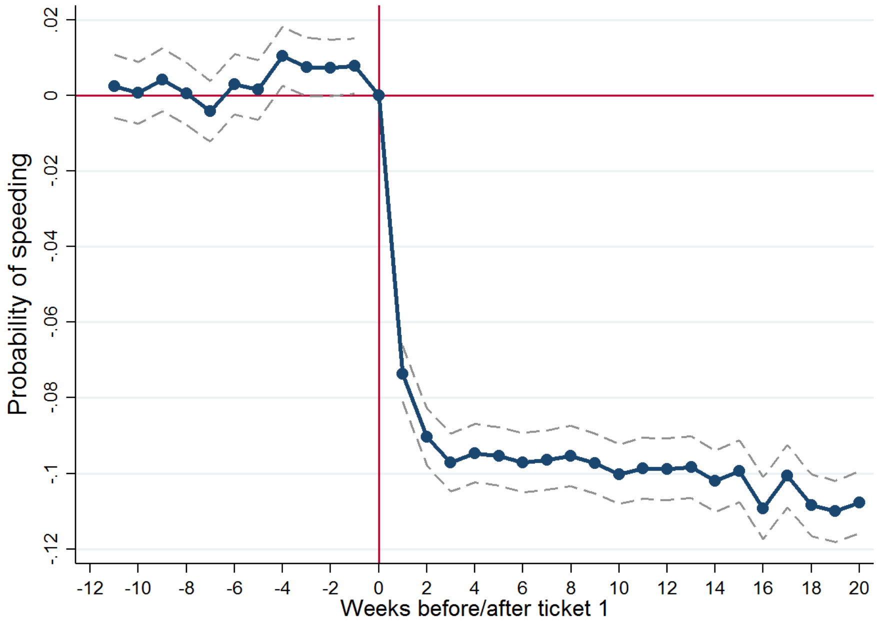

(b) Outcome: Speed

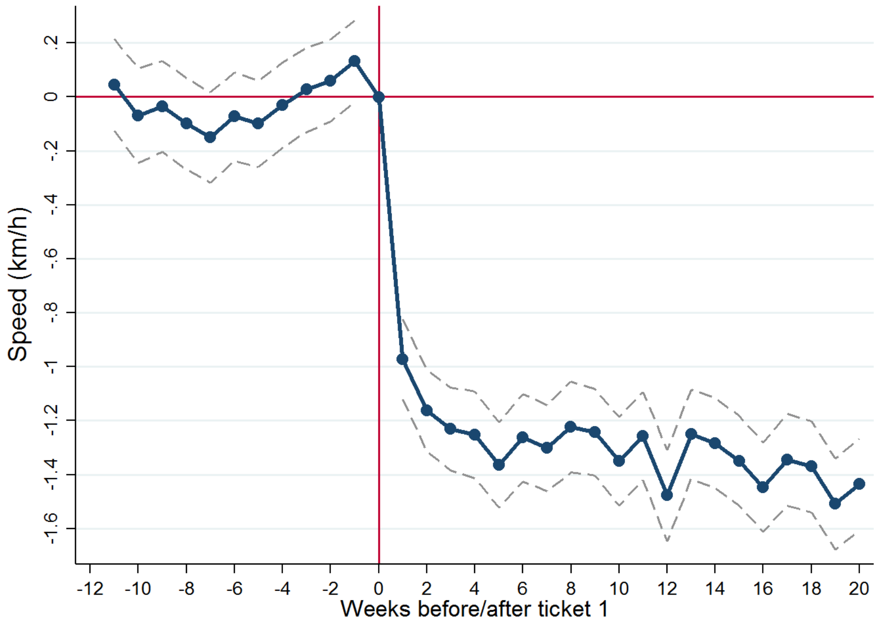

Notes: The figure plots the estimated β_w-coefficients from equ. (9) and their 95%-confidence intervals. Dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). (The corresponding estimates are also reported in Column (1) and (2) of Table A.11.) The sample includes cars around the first ticket event punished by the 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 (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: 0.27%.

Figure 11: Event study: shift in the speed distribution

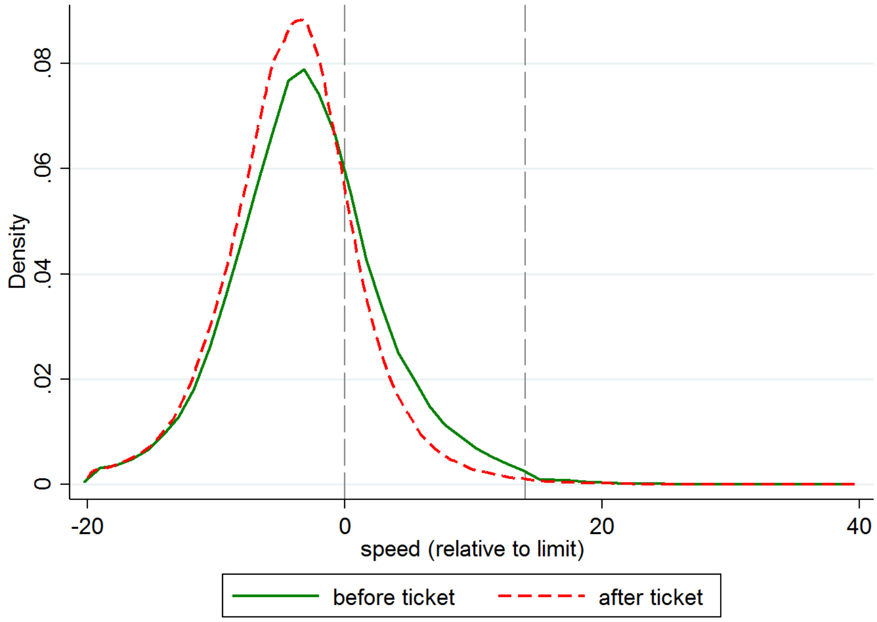

Notes: The figure depicts speed distributions in the event study sample (same as the sample used to generate estimates in Figure 10 and Table A.11). The speed is normalized relative to the speed limit. The solid green line plots the distribution for all rides made 12 and fewer weeks prior to receiving the 1st speeding ticket, with the trigger observation excluded. The dashed red line plots the distribution for all rides made during 20 weeks after receiving the ticket. The vertical lines mark the speed limit and the first enforcement cutoff.

Figure 12: Event study estimates: long-run effects

(a) Outcome: Speeding

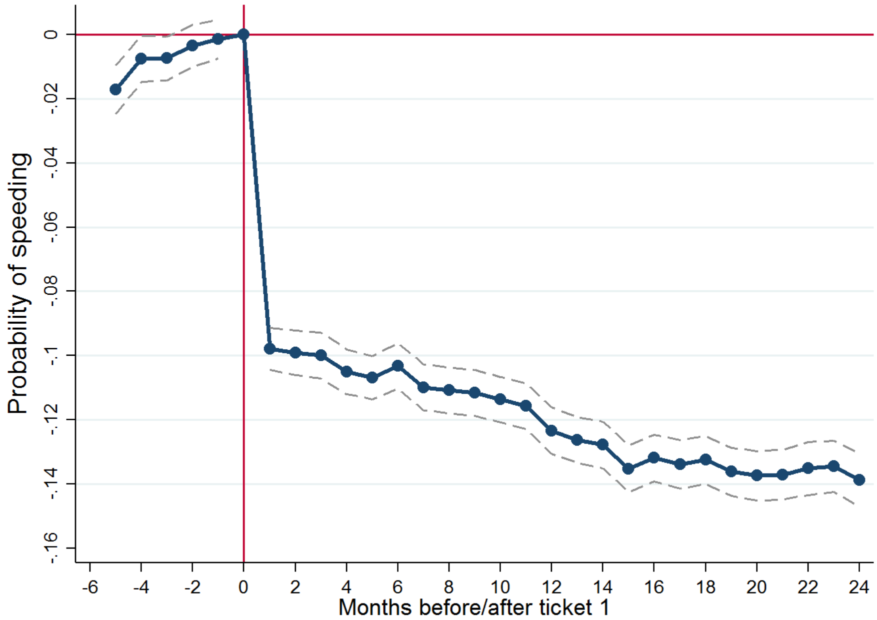

(b) Outcome: Speed

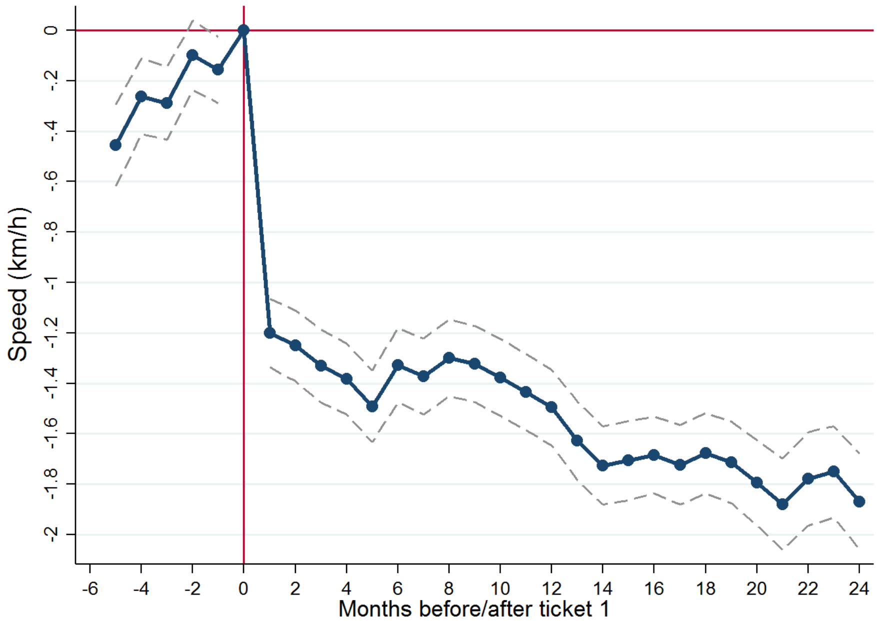

Notes: The figure plots the estimated β-coefficients and their 95%-confidence intervals from an equation analogous to equ. (9), where the single dummies indicate individual months (rather than weeks) before and after receiving a speeding ticket. Dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). (The corresponding estimates are also reported in Column (1) and (2) of Table A.12.) The sample includes cars around the first ticket event punished by the 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 (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: 0.26%.

Figure 13: Event study estimates: responses to the high- vs low-fine ticket

(a) Outcome: Speeding

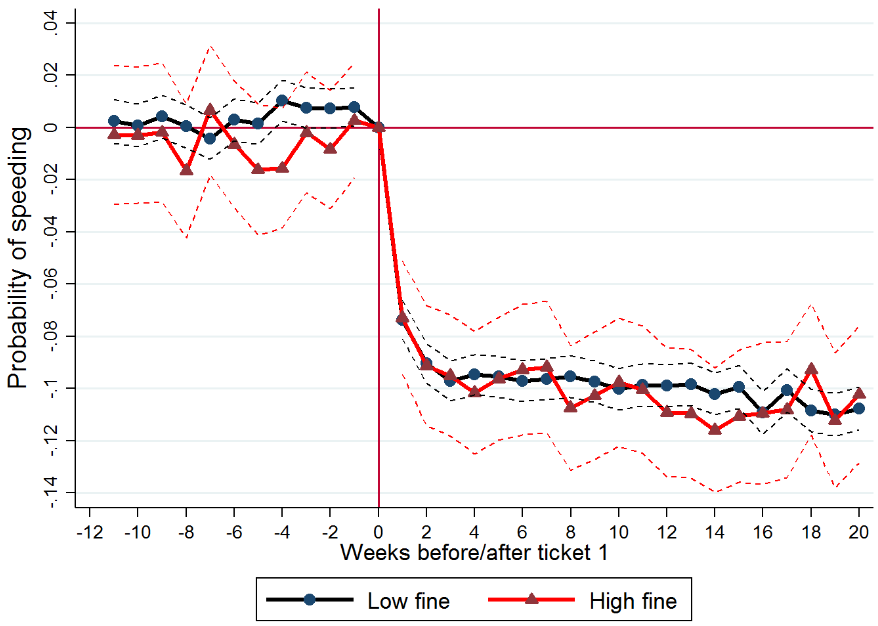

(b) Outcome: Speed

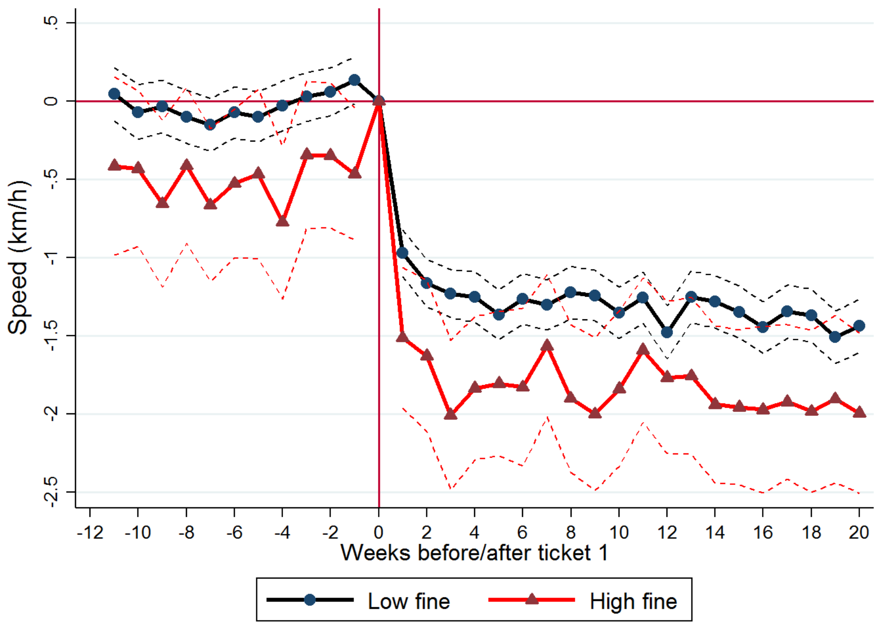

Notes: The figure plots the estimated β_w-coefficients from equ. (9) and their 95%-confidence intervals for cars receiving a high-fine tickets. The estimates for low-fine tickets (also displayed in Figure 10 above) are included for comparison. Dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). (The corresponding estimates are also reported in Table A.11.) The sample includes cars around the first ticket event punished by the 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. 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: 0.279%.

Figure 14: Event study estimates: private owner vs corporation

(a) Outcome: Speeding

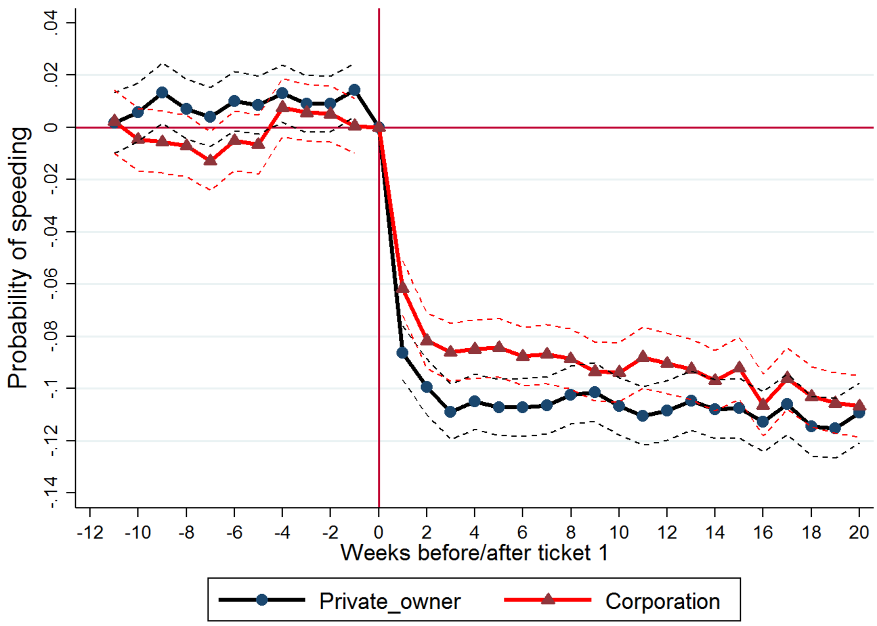

(b) Outcome: Speed

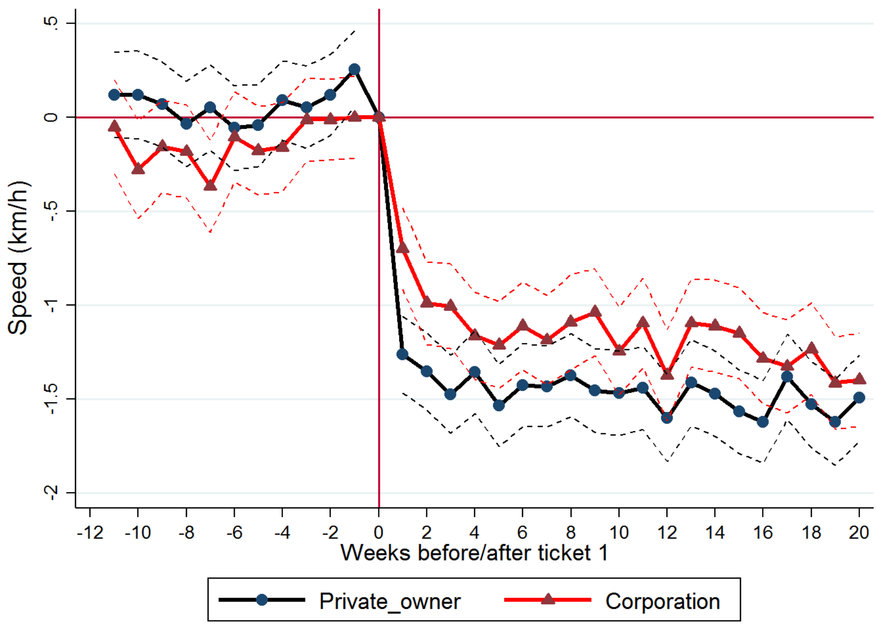

Notes: The figure plots the estimated β_w-coefficients from equ. (9) and their 95%-confidence intervals, separately for private and corporation cars. Dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). (The corresponding estimates are also reported in Table A.13.)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

Figure 15: Event study estimates: paid vs unpaid tickets

(a) Outcome: Speeding

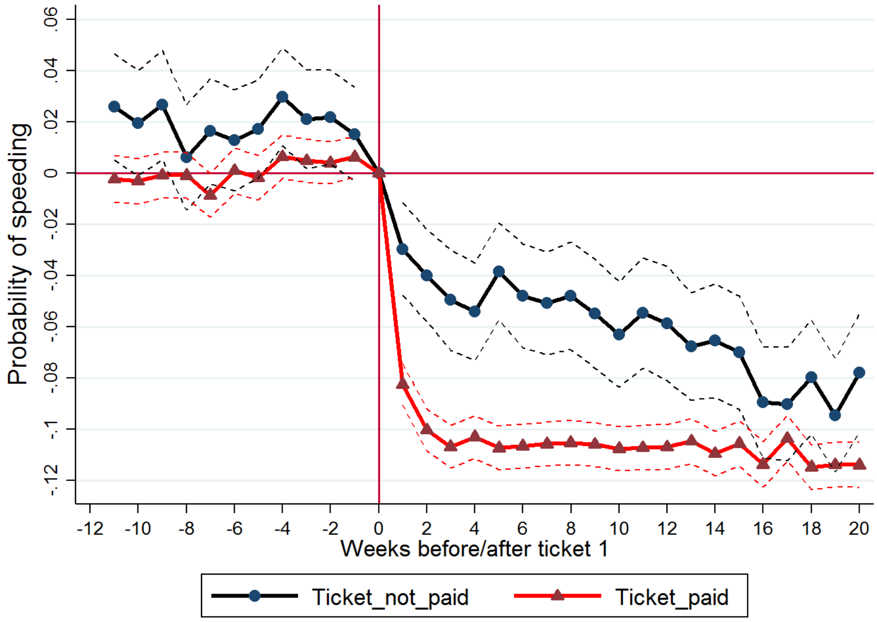

(b) Outcome: Speed

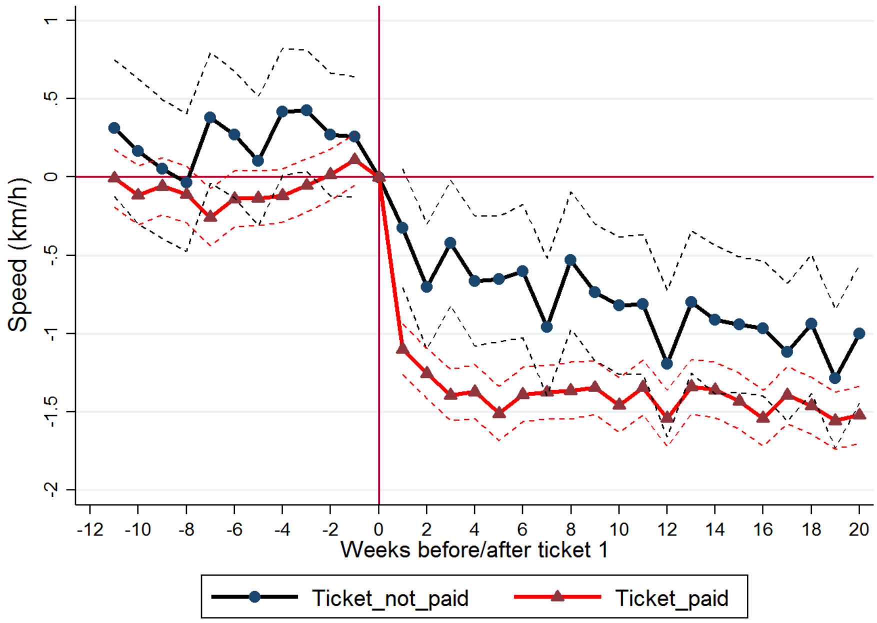

Notes: The figure plots the estimated β_w-coefficients from equ. (9) and their 95%-confidence intervals, separately for cars that paid the ticket within 90 days of receiving it and cars that did not. Dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). (The corresponding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

Figure 16: Event study estimates by the ride sequence

(a) Outcome: Speeding

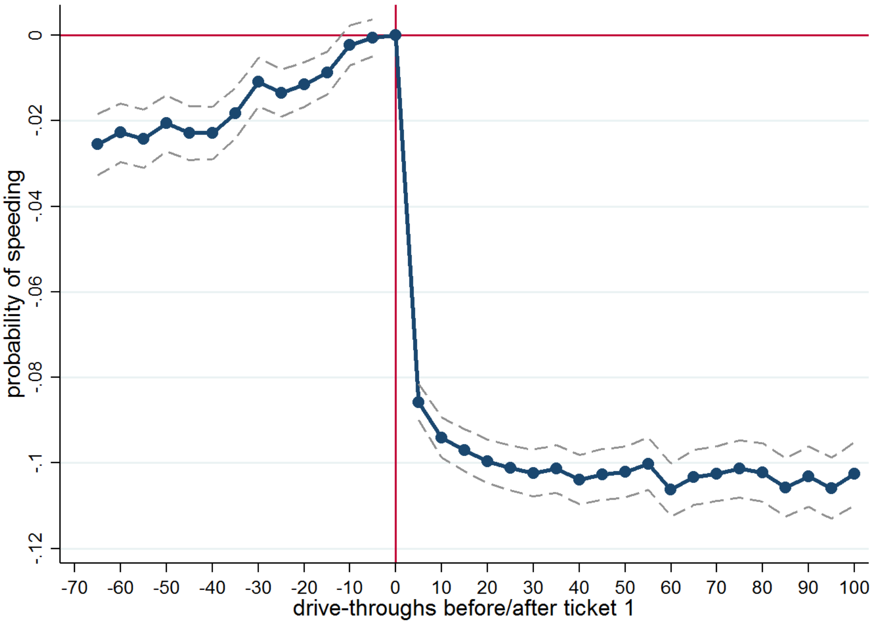

(b) Outcome: Speed

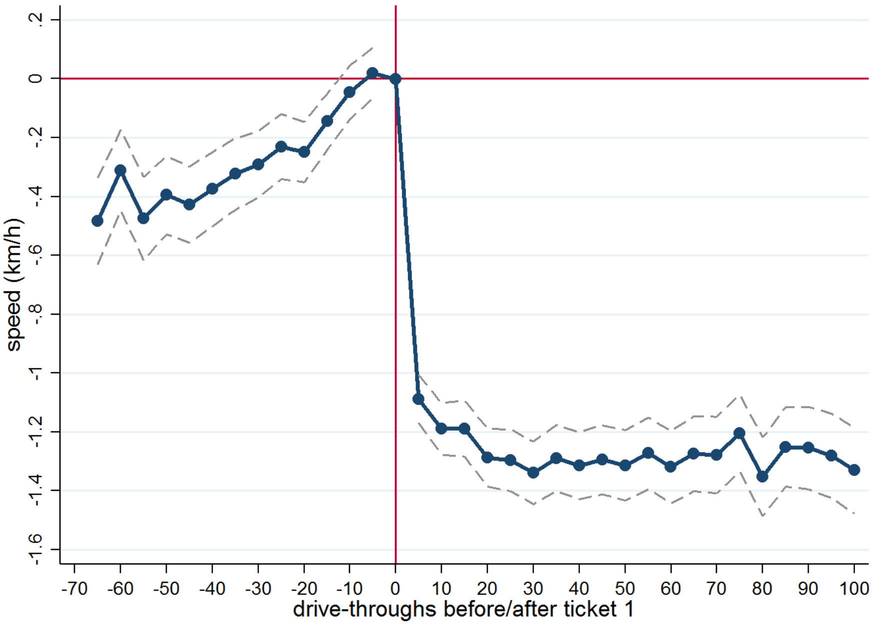

Notes: The figure plots the coefficients (and their 95% confidence intervals) from a regression is analogous to equ. (9) except that week dummies are replaced with indicators for a cars' ride sequence. Each dummy indicates five rides in their order before and after receiving the first ticket. Camera-specific calendar month fixed effects are also included. The sample of cars is identical to the main estimates presented in Figure 10. Dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). (The corresponding estimates are also reported in Table A.15.) 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) / 0.29%.

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

(a) Outcome: Speeding

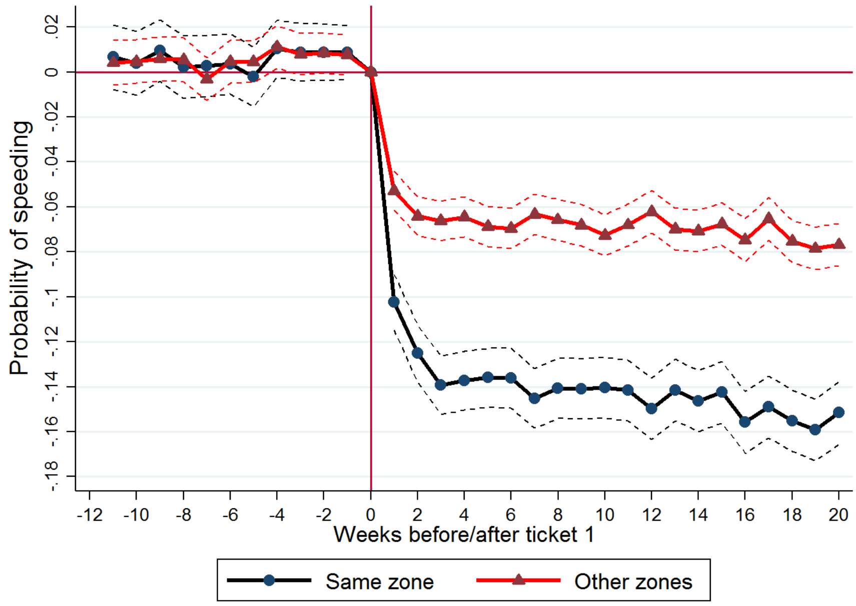

(b) Outcome: Speed

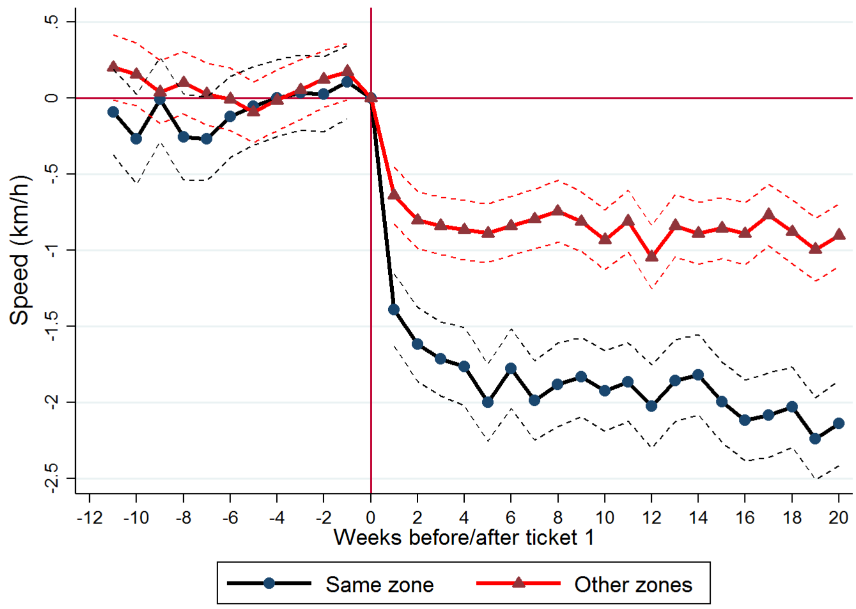

Notes: The figure plots the estimated β_w-coefficients from equ. (9) and their 95%-confidence intervals, separately for observations occurring at the same camera where the ticket was triggered and at the other speed cameras. Dependent variables are the binary speeding indicator (Panel a) and the mean speed (Panel b). The sample includes cars around the first ticket event punished by the 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 (last week before receiving the ticket) is the omitted category. The estimates are reported in Table A.16.

Figure 18: Event study estimates for responses to the second ticket

(a) Outcome: Speeding

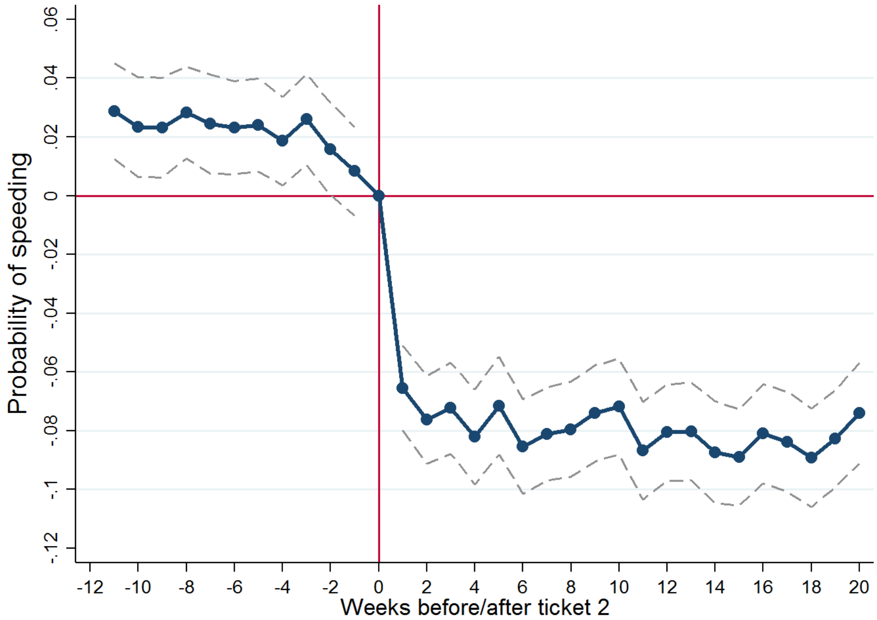

(b) Outcome: Speed

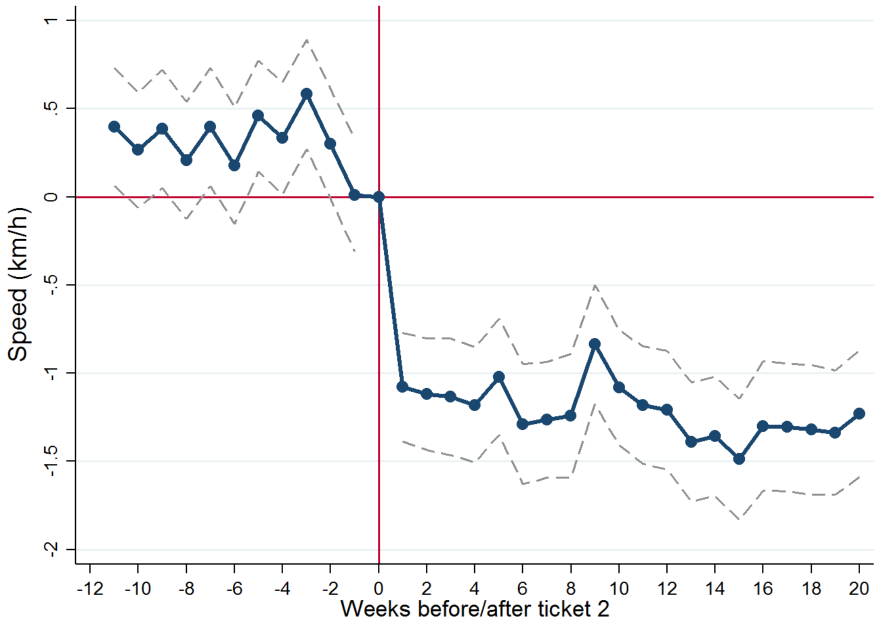

Notes: The figure plots the estimated β_w-coefficients from equ. (9) and their 95%-confidence intervals. Dependent variables are the mean speed (Panel a) and the binary speeding indicator (Panel b). The sample includes cars around the second ticket event (which previously experienced a first, low-fine 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 is available. The trigger observation is excluded. Week zero (last week before receiving the second ticket) is the omitted category. Cars: 2,566. Observations: 157,098. Mean speed/mean speeding rate in week zero: 45.03km/h / 0.27%. The estimates are also reported in Table A.11.

Tables

Table 1: Summary statistics

	'Not-ticketed' cars	'Ticketed' cars	Total (all cars)
Car characteristics			
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	84.36	19.31
	(74.84)	(192.02)	(82.93)
Driving frequency	2.33	3.06	2.45
	(2.76)	(2.87)	(2.79)
Number plate: Local region	0.453	0.455	0.453
	(0.498)	(0.498)	(0.498)
Number plate: Prague	0.393	0.439	0.400
	(0.488)	(0.496)	(0.490)
Ride characteristics			
Speed	–6.00	–5.17	–5.87
	(7.73)	(8.60)	(7.88)
Speeding	0.125	0.189	0.135
	(0.331)	(0.391)	(0.342)
(Re)offending	0.000	0.015	0.003
	–	(0.120)	(0.051)
Temperature	15.36	14.81	15.27
	(12.19)	(12.09)	(12.17)
Windspeed	1.75	1.74	1.75
	(1.46)	(1.47)	(1.46)
Hour	12.51	12.65	12.53
	(4.59)	(4.73)	(4.61)
Weekend	0.204	0.205	0.204
	(0.403)	(0.404)	(0.403)
Ticket/trigger characteristics			
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 parenthesis) for cars that did ('ticketed') or did not get any speeding ticket during the sample period (August 2014–2018). Speed indicates the measured speed, relative to the speed limit (in km/h). Number plate distinguish cars registered in the local region (Central Bohemian, 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/1st cutoff (enforcement)	yes (correct prior) no (incorrect prior)	yes (evolving over time)	no
Bunching/2nd 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 drop in speed$^{(b)}$ (if higher fine induces stronger updating)	no differential effect$^{(b)}$ (potentially larger drop in speed for favourable driving conditions)

Notes: (a) These prediction implicitly assumes 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. For the case of coarse updating, we further assume that the second cutoff is not known. If it were known, we should see bunching (at the second cutoff) under coarse updating, too. (b) On these predictions, see the left and the right panel of Figure 2, respectively.

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	–0.0951***	–0.0051***	–1.4602***	–1.3097***	–1.4972***	–1.7723***
($\beta^{k=1}$)	[0.0136]	[0.0019]	[0.2774]	[0.2794]	[0.2663]	[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 a 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
		1st cutoff			*2nd cutoff*	
Estimate	–0.0707***	–0.0031***	–0.8804***	–0.0279	–0.0025	–0.8247
(β^{k})	[0.0139]	[0.0009]	[0.3191]	[0.0271]	[0.0034]	[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 a MSE-optimal bandwidth and cluster robust standard errors in brackets (Calonico et al., 2014, 2017). Number of observations indicate 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
			(A) *Outcome: Speeding*		
Estimate	–0.1186***	–0.0761***	–0.0953***	–0.0781***	–0.0924**
($\beta^{k=1}$)	[0.0244]	[0.0181]	[0.0226]	[0.0262]	[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
			(B) *Outcome: Mean Speed*		
Estimate	–1.6591***	–1.2295***	–1.0293**	–1.8673***	–1.6526***
($\beta^{k=1}$)	[0.4994]	[0.3192]	[0.4150]	[0.5316]	[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 a 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) Speeding	(2) (Re)Offending	(3) Speed	(4) Speed^{p50}	(5) Speed^{p75}	(6) Speed^{p90}
Estimate	–0.0243	–0.0058	–0.7225	–0.6508	–0.8824	–0.6883
($\beta^{k=2}$)	[0.0288]	[0.0104]	[0.7913]	[0.7782]	[0.7895]	[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 a 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: 16,148 cars.

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

	(1) Speeding (binary)	(2) Speed (mean)	(3) Speed^{p90}	(4) Speeding (binary)	(5) Speed (mean)	(6) Speed^{p90}
Panel A.	*Good Conditions*			*Bad Conditions*		
Estimate ($\beta^{k=2}$)	–0.0808*	–1.4711*	–2.0812**	–0.0075	–0.0809	–0.5930
	[0.0471]	[0.8681]	[1.0525]	[0.0330]	[0.7750]	[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.	*Good Conditions*			*Bad Conditions*		
Estimate ($\beta^{k=2}$)	–0.0796**	–0.8873	–1.0067	–0.0038	–0.3259	–0.8595
	[0.0393]	[0.6962]	[0.7688]	[0.0270]	[0.5453]	[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 a 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 under 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 5.84 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 under 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(left).

Table 8: Wald estimates: ‘same’ vs ‘other’ radar: **enforcement cutoff**

	(1) same	(2) other	(3) same	(4) other
		(A) *Outcome: Speeding*		
Estimate	–0.1166***	–0.0612***	–0.1065***	–0.0591***
($\beta^{k=1}$)	[0.0184]	[0.0172]	[0.0207]	[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) *Outcome: Mean Speed*		
Estimate	–1.8922***	–1.0746***	–1.8827***	–0.9951***
($\beta^{k=1}$)	[0.3716]	[0.2933]	[0.3469]	[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 a 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 at the *same* radar that triggered the assignment speed, S_i. Columns (2) and (4) explore outcomes from *other* radars, i.e., radars that differ 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’ radar; 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(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 ℓ = 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 a 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(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 ℓ = 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 a 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(left).

Appendix 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')

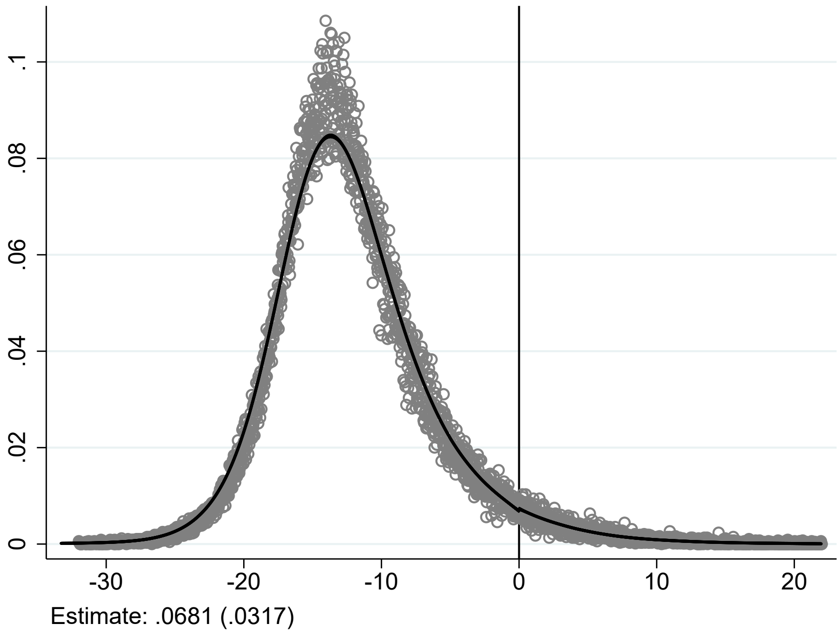

(b) Histogram

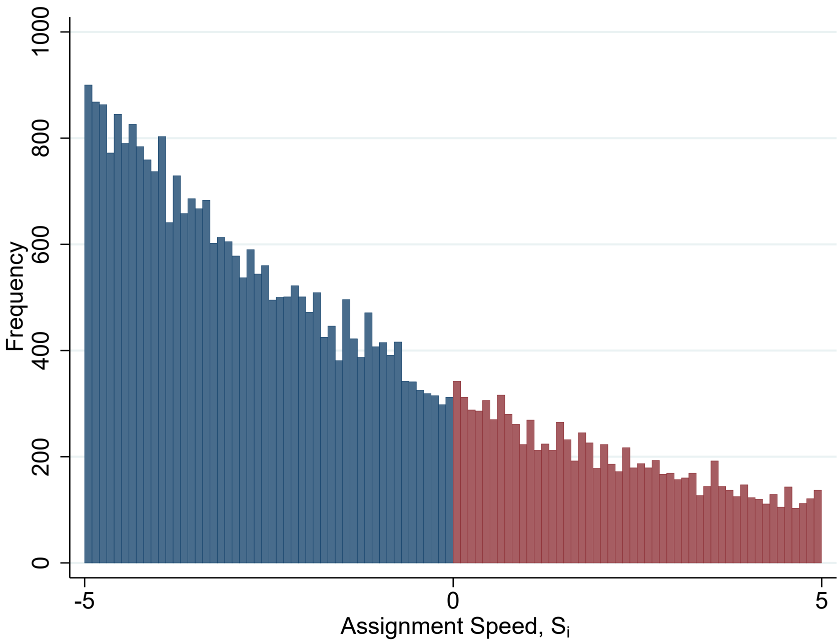

Notes: The figure illustrates the distribution of the assignment speed S_i centered around the enforcement cutoff (14km/h above the speed limit). Panel (a) plots the distribution together with the estimates from McCrary's (2008) heaping test. Panel (b) presents a histogram of the assignment speed S_i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

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

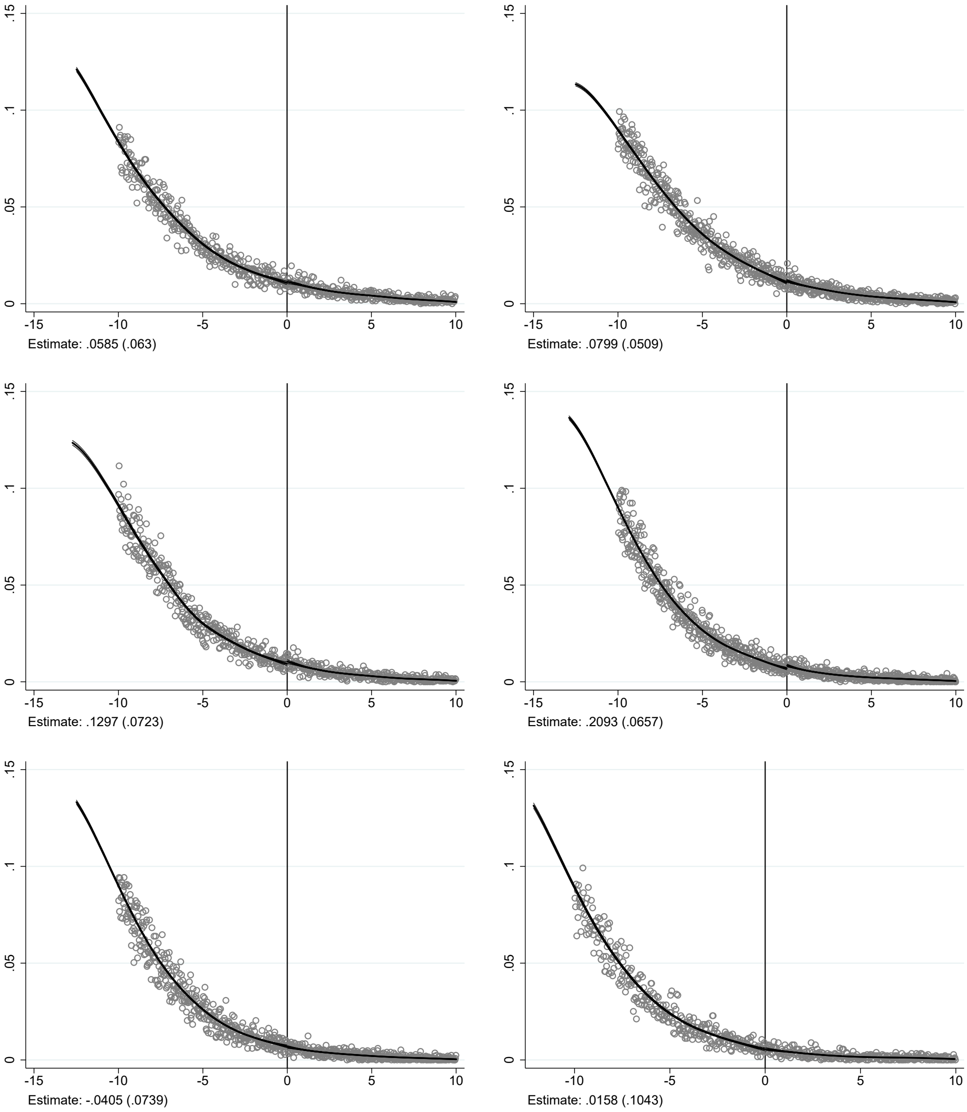

Notes: The figure plots the distribution of observed speed measures – centered around the enforcement cutoff (14km/h above the speed limit) – in six semi-annual intervals, starting with the first month the radars were operating. Sample includes all cars. Estimates are from McCrary's (2008) heaping test.

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

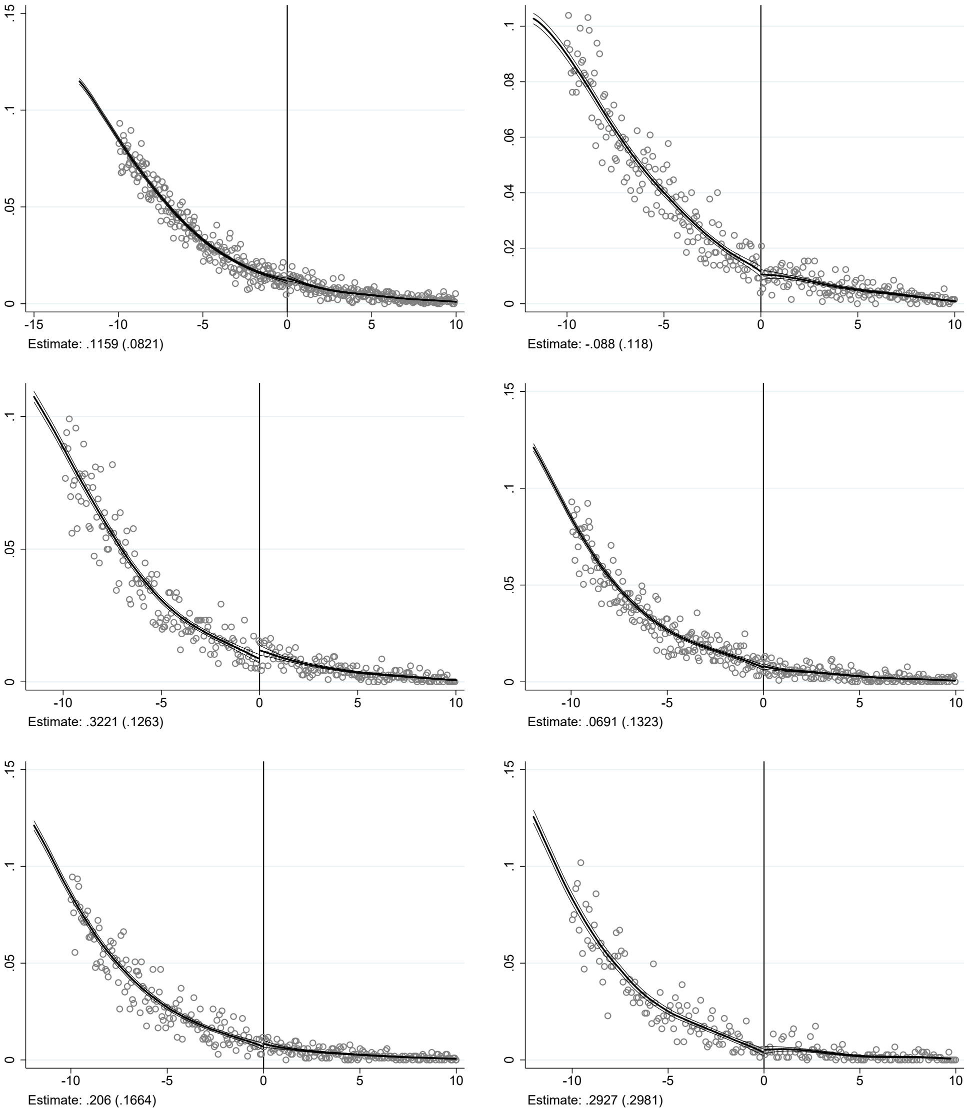

Notes: The figure plots the distribution of observe speed measures – centered around the enforcement cutoff (14km/h above the speed limit) – in six semi-annual intervals, starting with the first month the radars were operating. Sample includes only 'regional cars', defined as cars with above-median driving 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

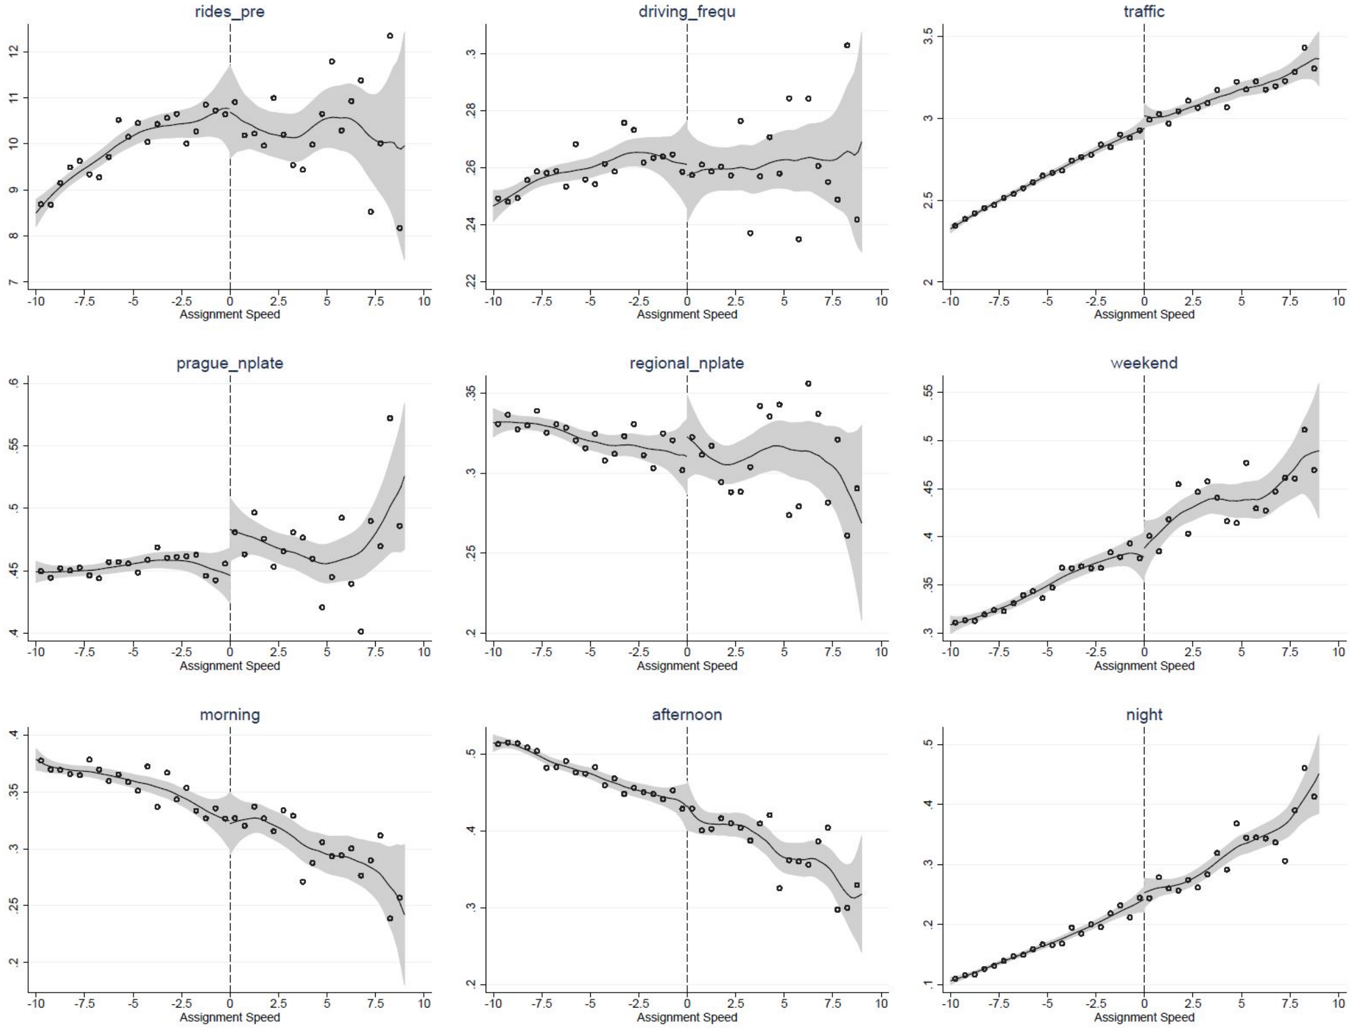

Notes: The figures document continuity for several observables around the enforcement cutoff. The figures depict the cars' number of rides as well as the driving frequency during the assignment period (pre-treatment); the traffic density (the logged time difference to the car ahead, measured for the 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')

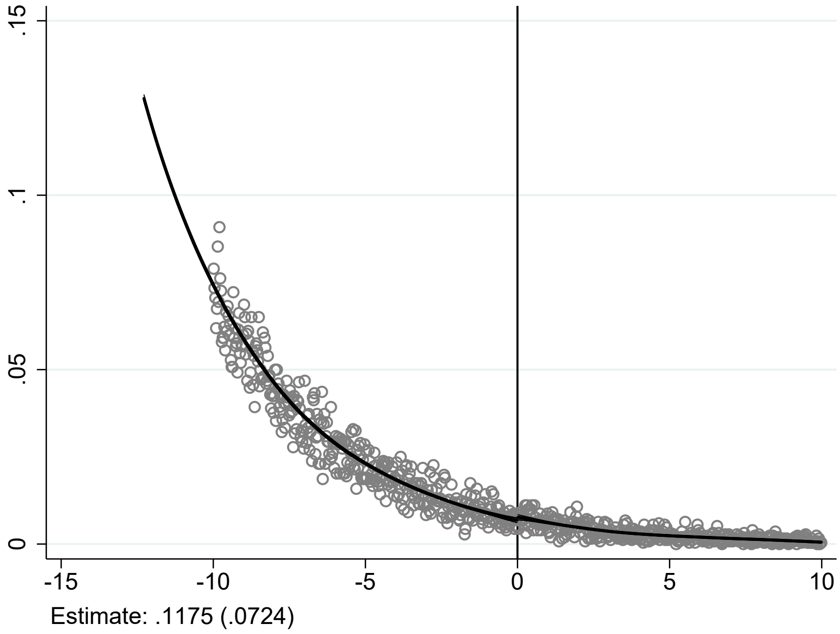

(b) Histogram

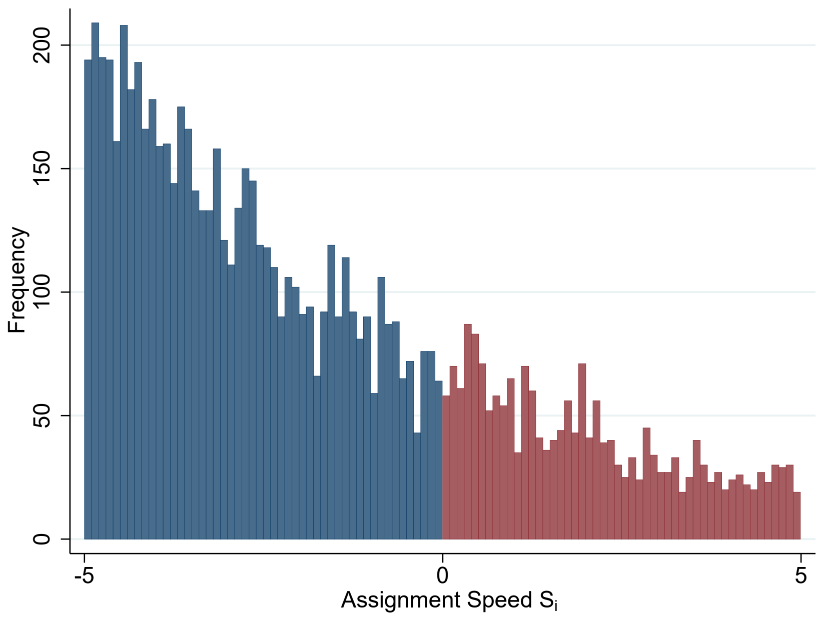

Notes: The figure illustrates the distribution of the assignment speed S_i centered around the high-fine cutoff (23km/h above the speed limit). Panel (a) plots the distribution together with the estimates from McCrary's (2008) heaping test. Panel (b) presents a histogram of the assignment speed S_i over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

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

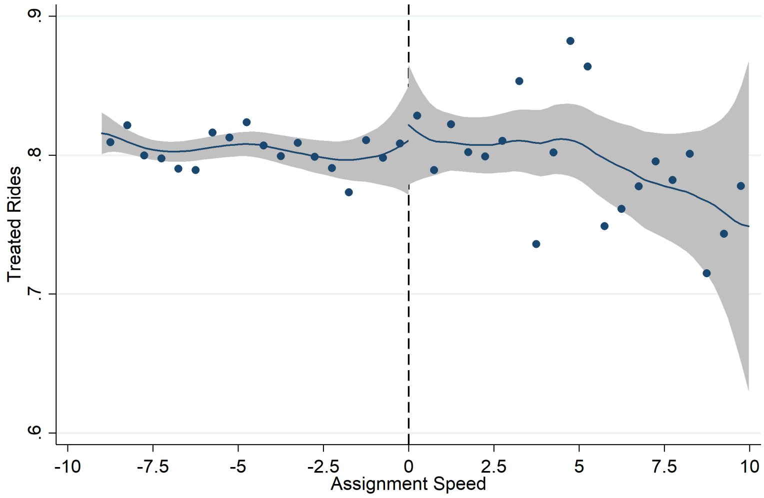

Notes: The figure presents the cars' share of *ticketed rides* T_i^1, i.e., rides after receiving any speeding ticket (relative to all rides in the outcome period), around the *high-fine cutoff* (2nd cutoff). If high-fine tickets are sent out more quickly than low-fine tickets, this could produce a discontinuity 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: Sensitivity of reduced form estimates at the enforcement cutoff

(a) Outcome: Speeding

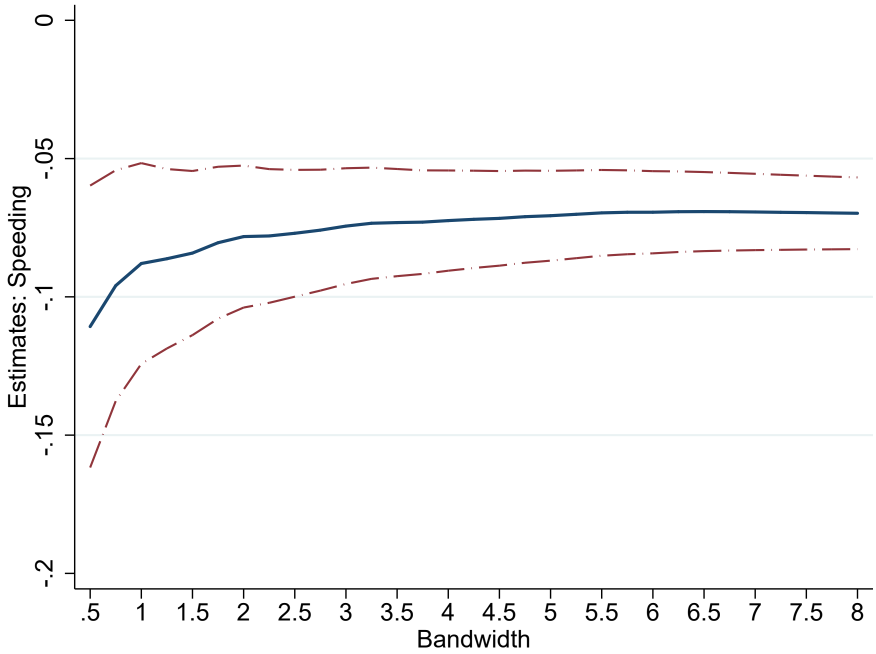

(b) Outcome: Speed

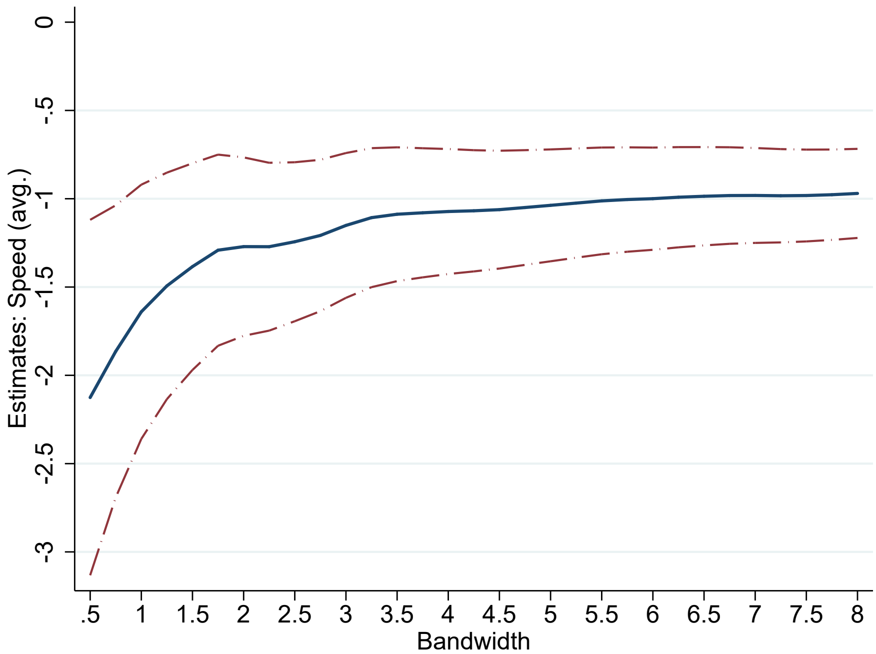

Notes: The figure depicts reduced form estimates (with 95% confidence intervals) for car-level observations at the enforcement cutoff, varying the bandwidth in 0.5km/h steps from 0.5 to 8.0km/h of assignment speed S_i. The outcome is the speeding (top panel) and the mean speed (lower panel). The different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

Figure A.9: Sensitivity of *car-level* estimates at the enforcement cutoff

(a) Outcome: Speeding

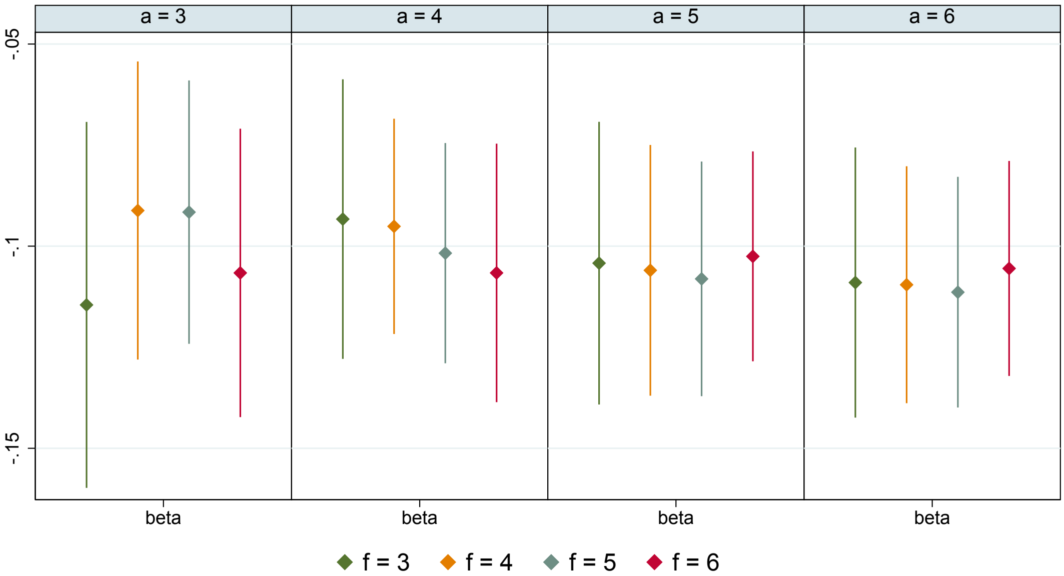

(b) Outcome: Speed

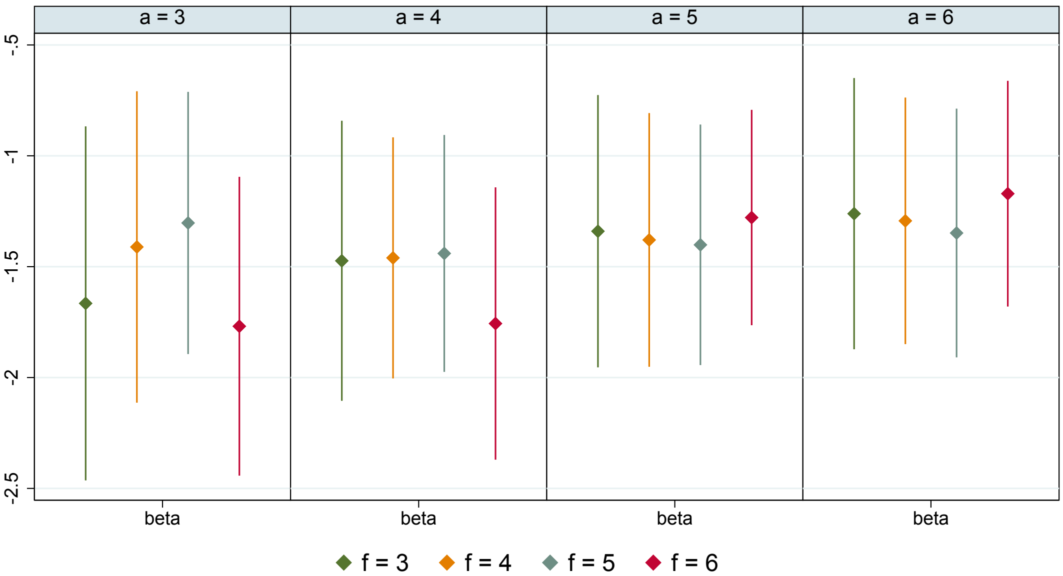

Notes: The figure depicts Wald estimates at the car-level, with 95% CI for the enforcement cutoff (1st cutoff) for different assignment (a, in months) and follow-up periods (f). Panel (a) presents outcomes for or the speeding rate, panel (b) the mean speed (in km/h). Number of observations, bandwidth,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

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

(a) Outcome: Speeding

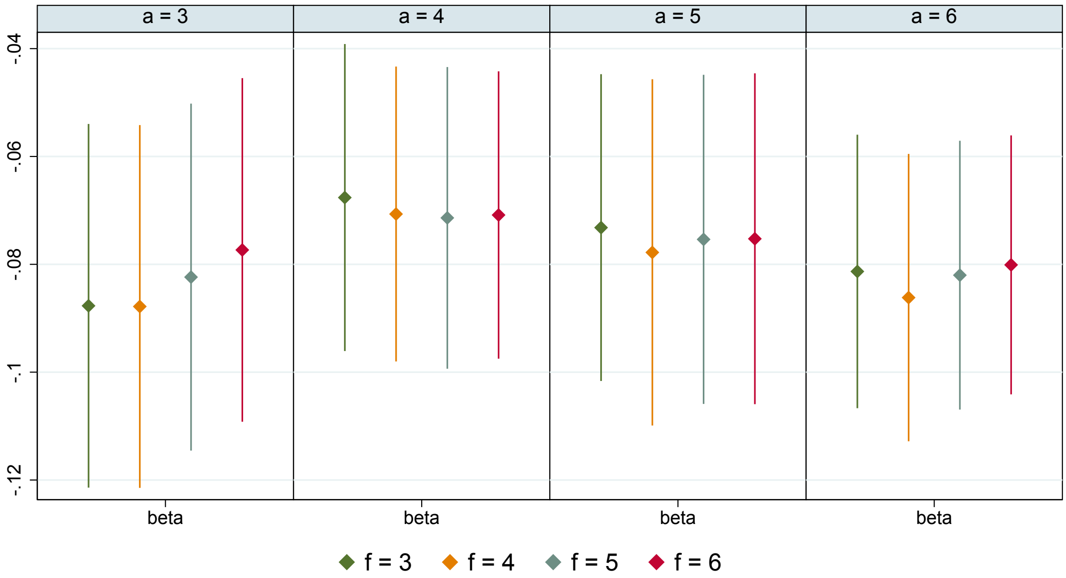

(b) Outcome: Speed

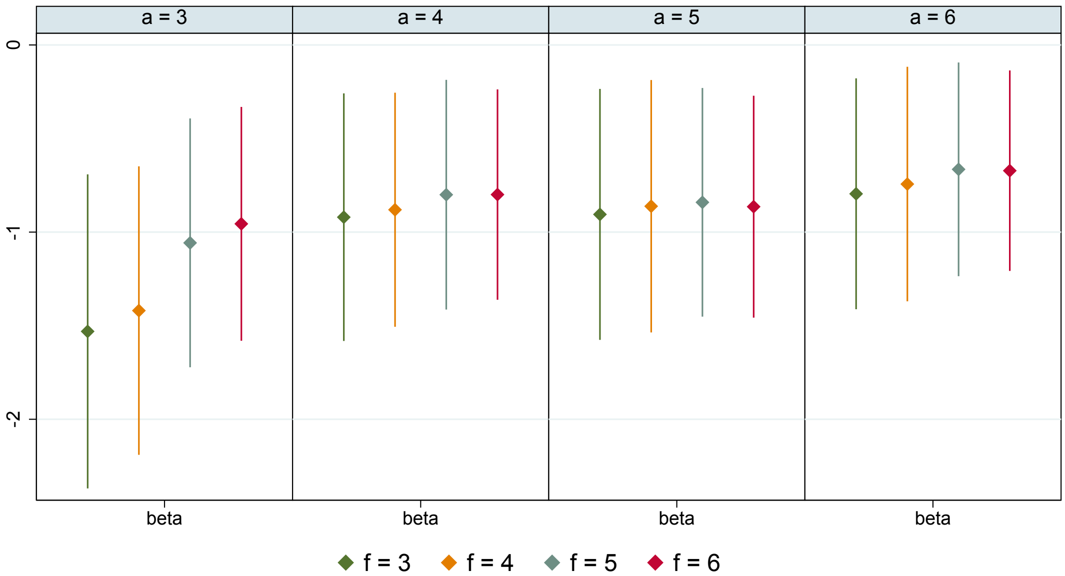

Notes: The figure depicts Wald estimates at the *ride-level*, with 95% CI for the enforcement cutoff (1st cutoff) for different assignment (a, in months) and follow-up periods (f). Panel (a) presents outcomes for the speeding rate, panel (b) for the mean speed (in km/h).

Figure A.11: Sensitivity of reduced form estimates at the high-fine cutoff

(a) Outcome: Speeding

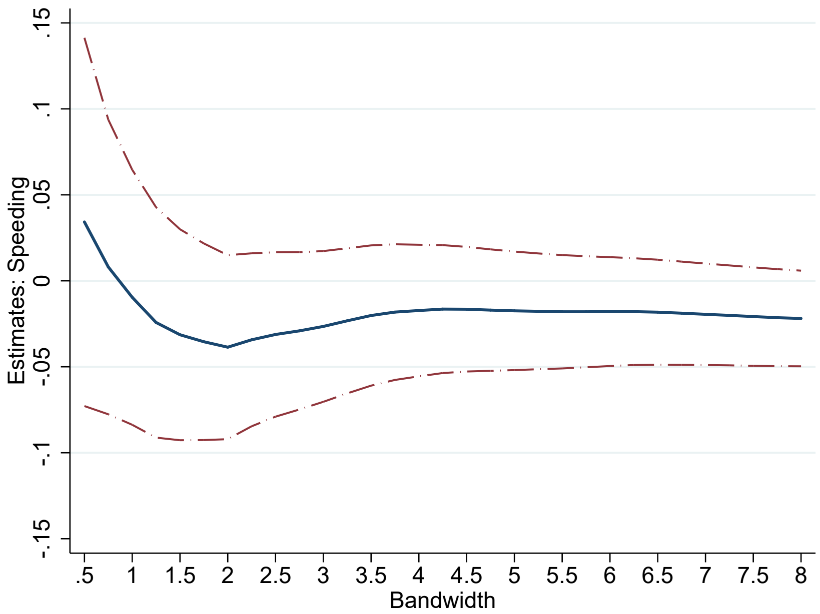

(b) Outcome: Speed

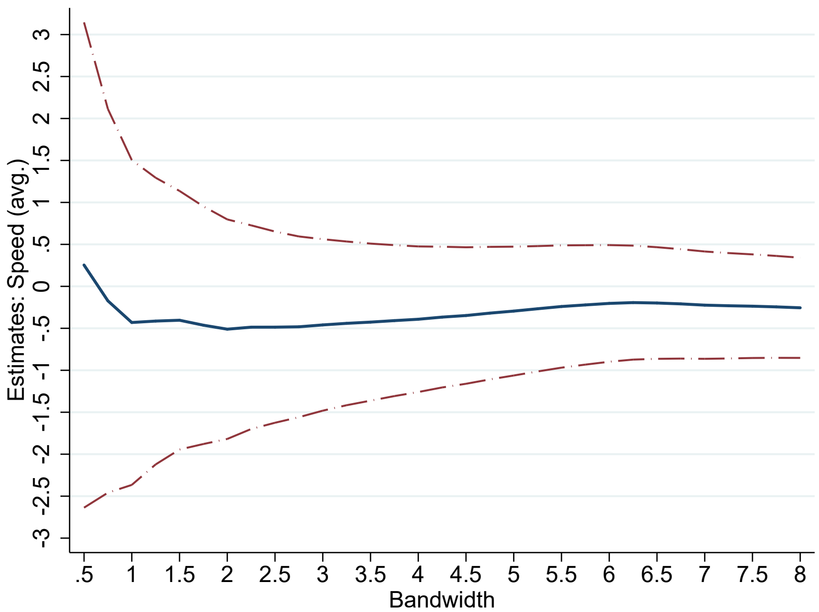

Notes: The figure depicts reduced form estimates (with 95% confidence intervals) for car-level observations at the high-fine cutoff, varying the bandwidth in 0.5km/h steps from 0.5 to 8.0km/h of assignment speed S_i. Outcome: Speed (in km/h). Panel (a) presents outcomes for the speeding rate, panel (b)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

Figure A.12: Sensitivity of *car-level* estimates at the high-fine cutoff

(a) Outcome: Speeding

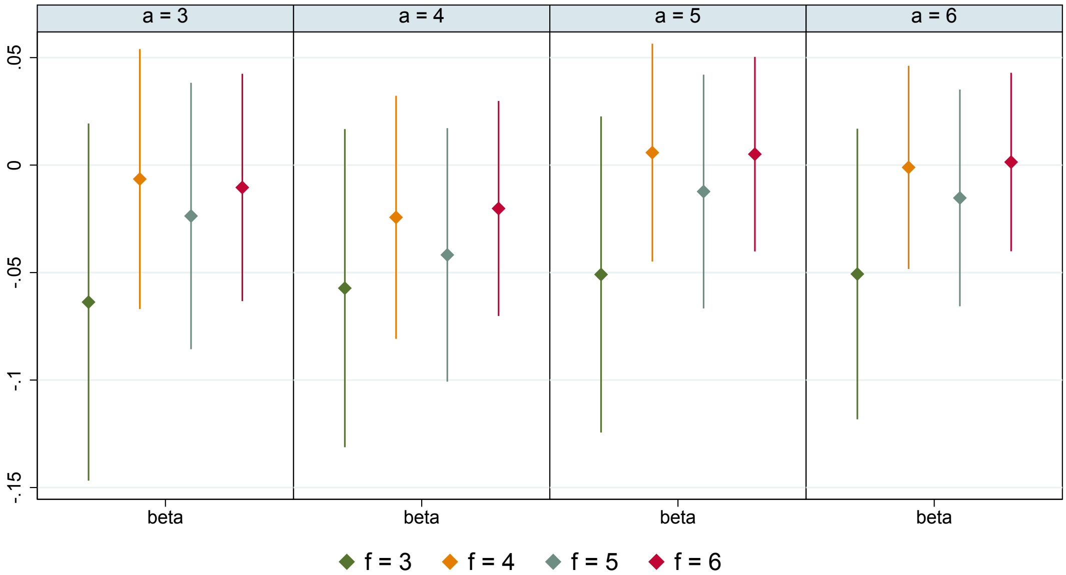

(b) Outcome: Speed

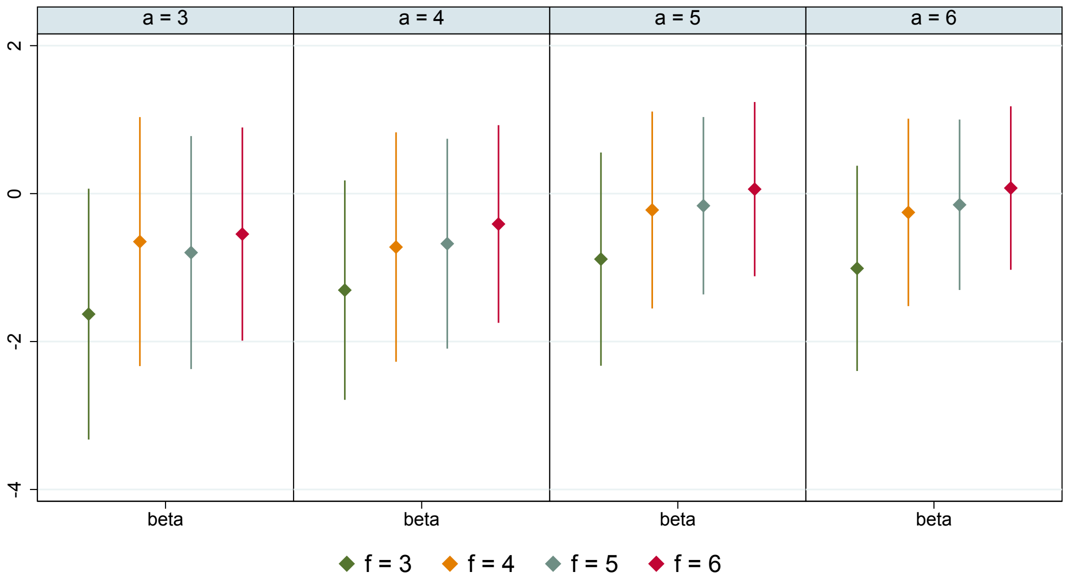

Notes: The figure depicts Wald estimates (at the car-level) with 95% CI for the high-fine cutoff (2nd cutoff) for different assignment (a, in months) and follow-up periods (f). Outcome: Speed (in km/h). Observations pooled at level of cars.

Figure A.13: Heterogeneity: speed and speeding responses by driving frequency

(a) Outcome: Speeding

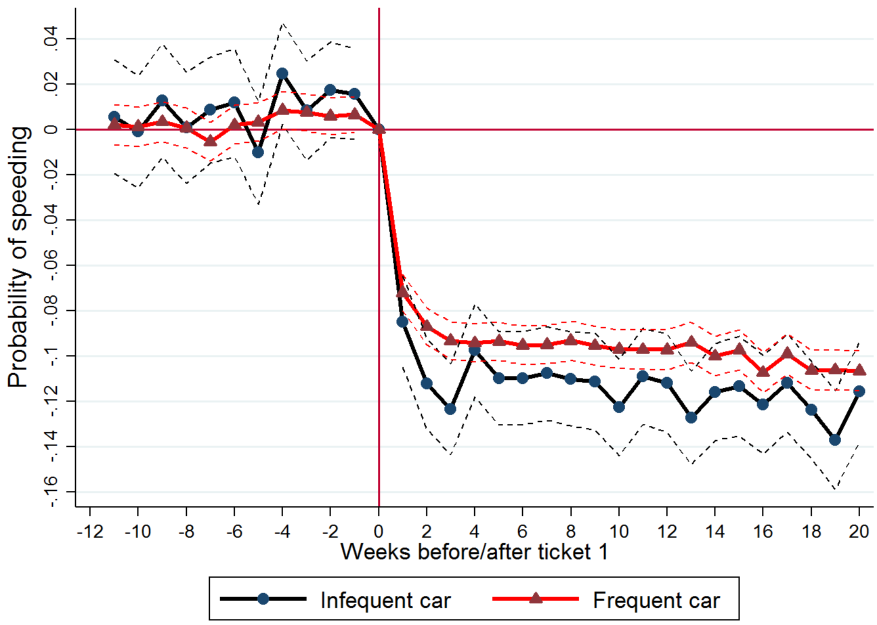

(b) Outcome: Speed

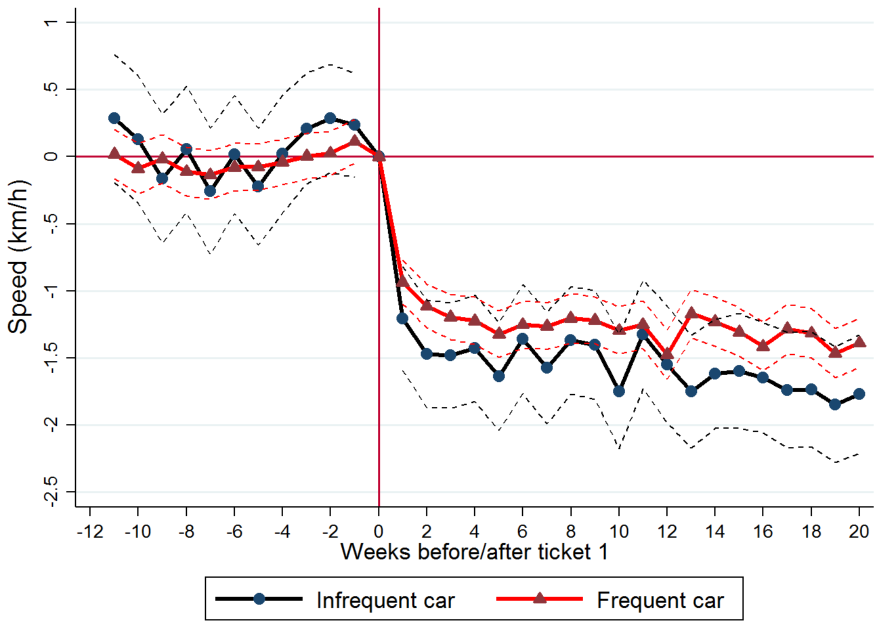

Notes: The figure plots the estimated β_w-coefficients from equ. (9) and their 95%-confidence intervals. Dependent variables are the speeding dummy (Panel a) and measured speed s_{it} in km/h(b). The sample is divided into infrequent and frequent cars, as defined by a median split according to the average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

Figure A.14: Heterogeneity: speed and speeding responses by number plate region

(a) Outcome: Speeding

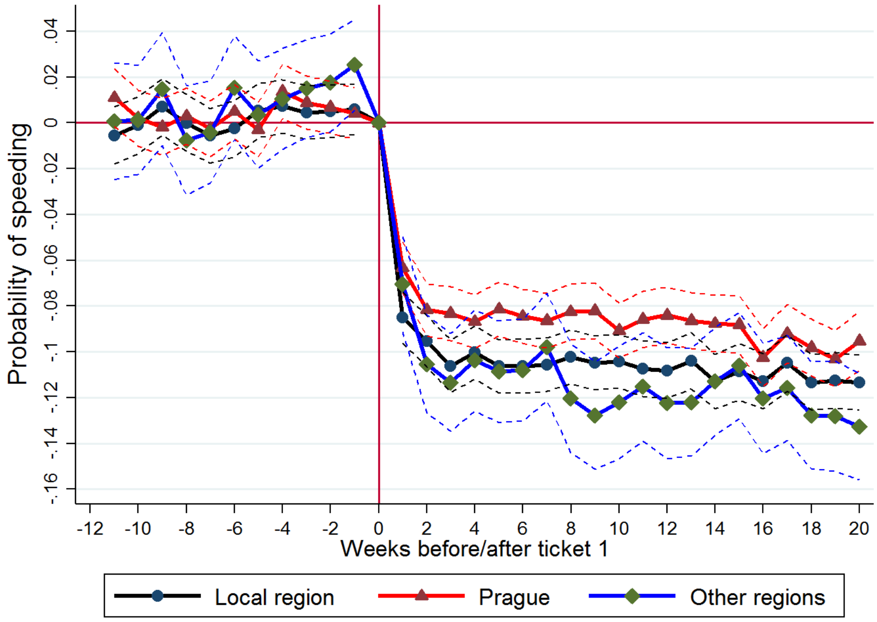

(b) Outcome: Speed

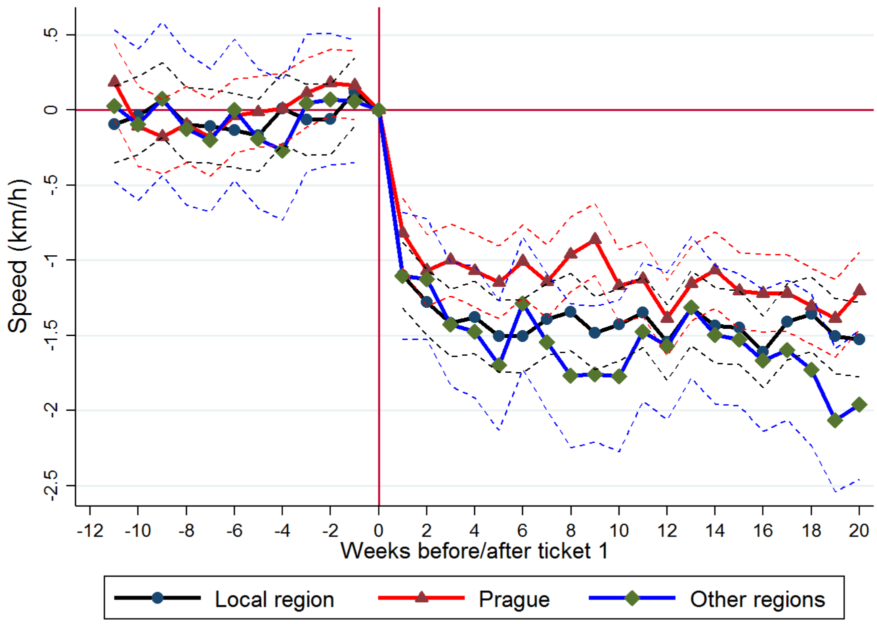

Notes: The figure plots the estimated β_w-coefficients from equ. (9) and their 95%-confidence intervals. Dependent variables are the speeding dummy (Panel a) and measured speed s_{it} in km/h (b). The sample is split by the number plate into cars from the 'Local' region (i.e., where the municipality 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 (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	–0.1127	0.0384*	0.0182	0.0915	–0.7626**	–0.1172	0.0211	–0.0306	0.0139
	[0.0135]	[0.8298]	[0.0204]	[0.0206]	[0.0587]	[0.3722]	[0.0819]	[0.0167]	[0.0220]	[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) 3–6pm	6–9pm	9–12am	day of week (trigger ride) Tue	Wed	Thu	Fr	Sa	Su	month 02
Estimate	0.0108	0.0167	0.0067	–0.0173	0.0081	0.0215	0.0058	0.0122	–0.0029	0.0192*
	[0.0195]	[0.0167]	[0.0114]	[0.0175]	[0.0153]	[0.0171]	[0.0140]	[0.0166]	[0.0187]	[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	0.0151	–0.0173	–0.0319**	0.0098	0.0018	0.0192*	–0.0039	–0.0032	–0.0087
	[0.0145]	[0.0137]	[0.0152]	[0.0132]	[0.0144]	[0.0114]	[0.0101]	[0.0113]	[0.0109]	[0.0151]
Y(left)	0.098	0.071	0.103	0.124	0.087	0.065	0.042	0.080	0.083	0.100

Notes: The table presents a series of balancing checks for the enforcement cutoff. In particular, the table reports bias-corrected RD estimates with a 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 (Region and Prague; residual category: all Others), as well as driving conditions (temperature, wind, and traffic density, measured via the logged time difference to the next car in front) and time and date information for the 'trigger ride' (ride the maximum speed). Y(left) indicates the mean of the dependent variable in the 0.5km/h bin below the cutoff. Number of observations for all specifications: 224,816 cars.

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	1.0433	–0.0151	0.0792*	0.0499	–0.1306	0.0948	0.0413	0.0278	0.0106
	[0.0230]	[1.5120]	[0.0422]	[0.0450]	[0.0993]	[0.9319]	[0.1520]	[0.0375]	[0.0334]	[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	0.0148	–0.0070	0.0000	0.0070	–0.0246	0.0274	–0.0574	0.0170	0.0422**
	[0.0355]	[0.0321]	[0.0264]	[0.0264]	[0.0321]	[0.0333]	[0.0260]	[0.0412]	[0.0447]	[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	0.0106	0.0105	–0.0659**	–0.0035	–0.0358	0.0428**	–0.0174	–0.0200	–0.0030
	[0.0226]	[0.0231]	[0.0278]	[0.0307]	[0.0311]	[0.0258]	[0.0210]	[0.0242]	[0.0210]	[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: The table presents a series of balancing checks for the high-fine cutoff. In particular, the table reports bias-corrected RD estimates with a 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 (Region and Prague; residual category: all Others), as well as driving conditions (temperature, wind, and traffic density, measured via the logged time difference to the next car in front) and time and date information for the 'trigger ride' (ride the maximum speed). Y(left) indicates the mean of the dependent variable in the 0.5km/h bin below the cutoff. Number of observations for all specifications: 16,148 cars.

Table A.3: Reduced form estimates: enforcement cutoff

	(1) Ticketed	(2) Speeding	(3) (Re)Offending	(4) Speed	(5) Speed^{p90}
Estimate	0.7866***	–0.0812***	–0.0044**	–1.3512***	–1.4023***
(δ, τ)	[0.0127]	[0.0146]	[0.0019]	[0.2814]	[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, in particular, bias-corrected RD estimates with a 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 (2nd cutoff)	(3) Speeding	(4) (Re)Offending	(5) Speed	(6) Speed^{p90}
Estimate	0.8145***	0.0117	–0.0199	–0.0048	–0.5938	–0.5632
(δ, τ)	[0.0189]	[0.0350]	[0.0235]	[0.0085]	[0.6494]	[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, in particular, bias-corrected RD estimates with a 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)
	1st cutoff		*2nd cutoff*	
Estimate	0.8812	0.0389**	–0.2831	0.0022
(τ)	[0.6501]	[0.0173]	[1.6193]	[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 cutoff (Columns 3 – 4), respectively. The sample is defined as in our main estimates (see, e.g., 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 is ever observed again (Columns 2 and 4). Bias-corrected RD estimates based on car-level observations with a 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)	a=3	-1.6654***	-1.4110***	-1.3028***	-1.7686***
		[0.4073]	[0.3583]	[0.3017]	[0.3438]
Y(left)		46.273	46.205	46.238	46.287
Relative effect		-3.599	-3.054	-2.818	-3.821
Bandwidth		2.580	3.042	3.763	2.439
Obs		203,459	221,303	235,072	246,178
Estimate (β)	a=4	-1.4734***	-1.4602***	-1.4402***	-1.7563***
		[0.3222]	[0.2774]	[0.2726]	[0.3133]
Y(left)		46.225	46.153	46.204	46.243
Relative effect		-3.187	-3.164	-3.117	-3.798
Bandwidth		3.423	4.199	3.774	2.515
Obs		207,737	224,816	239,220	250,887
Estimate (β)	a=5	-1.3400***	-1.3793***	-1.4013***	-1.2784***
		[0.3133]	[0.2919]	[0.2768]	[0.2478]
Y(left)		46.148	46.083	46.121	46.160
Relative effect		-2.904	-2.993	-3.038	-2.769
Bandwidth		3.017	3.150	3.074	3.436
Obs		211,031	227,732	241,375	253,579
Estimate (β)	a=6	-1.2608***	-1.2932***	-1.3481***	-1.1707***
		[0.3120]	[0.2837]	[0.2862]	[0.2597]
Y(left)		46.033	46.006	46.000	45.967
Relative effect		-2.739	-2.811	-2.931	-2.547
Bandwidth		2.915	3.136	2.670	2.913
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: 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(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.9.

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***	-0.0912***	-0.0916***	-0.1066***
		[0.0231]	[0.0188]	[0.0166]	[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***	-0.0951***	-0.1017***	-0.1066***
		[0.0176]	[0.0136]	[0.0139]	[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***	-0.1060***	-0.1081***	-0.1025***
		[0.0178]	[0.0158]	[0.0148]	[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***	-0.1095***	-0.1114***	-0.1055***
		[0.0170]	[0.0150]	[0.0146]	[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(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.9.

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.	*Good Conditions*			*Bad Conditions*		
Estimate	–0.1512***	–2.4135***	–2.7780***	–0.0536***	–0.7510***	–0.7596**
($\beta^{k=1}$)	[0.0218]	[0.4252]	[0.4723]	[0.0133]	[0.2765]	[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.	*Good Conditions*			*Bad Conditions*		
Estimate	–0.1283***	–1.8548***	–2.4969***	–0.0538***	–0.8609***	–0.8842***
($\beta^{k=1}$)	[0.0149]	[0.3454]	[0.3902]	[0.0130]	[0.2835]	[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 a 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 under 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 under 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 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(left).

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

	(1) Infrequent	(2) Frequent	(3) Local region	(4) Prague	(5) Other regions
	(A) *Outcome: Speeding*				
Estimate	–0.0510	–0.0113	–0.0732	0.0229	–0.0644
($\beta^{k=2}$)	[0.0432]	[0.0425]	[0.0480]	[0.0451]	[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) *Outcome: Mean Speed*				
Estimate	–0.9985	–0.5009	0.3482	–0.5939	–2.1516
($\beta^{k=2}$)	[0.9547]	[1.2346]	[1.2142]	[0.9385]	[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 a 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 A.10: Reduced form estimates: driving outcomes on unmonitored road (1st cutoff)

	car-level		*ride-level*	
	(1)	(2)	(3)	(4)
	Speed	Speed^{p90}	Speed	log(time)
Estimate	–1.1174	-3.9075*	–0.6468	0.0645**
(τ)	[1.8909]	[2.0315]	[0.4417]	[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 at the un-monitored road (the stretch of road between existing speed cameras No. 1 and entering into camera zone no. 4). The table reports both, estimates at the level of cars (Columns 1–2). and at the level of rides (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. Similar 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 a MSE-optimal bandwidth and robust (Columns 1–2) and clustered (3–4) standard errors in brackets (Calonico et al., 2014, 2017). Y(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: low- and high-fine tickets

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Low Fine Ticket				*High Fine Ticket*			
	first ticket		second ticket		first ticket		second ticket	
	Speeding	Speed	Speeding	Speed	Speeding	Speed	Speeding	Speed
Week -11 before	0.002	0.045	0.029***	0.397**	-0.003	-0.415	0.042	1.028*
	(0.004)	(0.087)	(0.008)	(0.171)	(0.014)	(0.290)	(0.026)	(0.549)
Week -10 before	0.001	-0.07	0.023***	0.265	-0.003	-0.431*	0.009	1.021*
	(0.004)	(0.089)	(0.009)	(0.167)	(0.013)	(0.254)	(0.025)	(0.558)
Week -9 before	0.004	-0.034	0.023***	0.386**	-0.002	-0.656**	0.013	0.465
	(0.004)	(0.086)	(0.009)	(0.171)	(0.014)	(0.272)	(0.023)	(0.558)
Week -8 before	0.001	-0.099	0.028***	0.207	-0.016	-0.411	0.006	0.516
	(0.004)	(0.086)	(0.008)	(0.169)	(0.013)	(0.254)	(0.026)	(0.669)
Week -7 before	-0.004	-0.151*	0.024***	0.397**	0.007	-0.664***	0.011	0.624
	(0.004)	(0.086)	(0.009)	(0.171)	(0.013)	(0.249)	(0.025)	(0.557)
Week -6 before	0.003	-0.073	0.023***	0.178	-0.006	-0.525**	0	0.675
	(0.004)	(0.084)	(0.008)	(0.169)	(0.012)	(0.243)	(0.026)	(0.561)
Week -5 before	0.002	-0.1	0.024***	0.458***	-0.016	-0.466*	0.013	0.508
	(0.004)	(0.082)	(0.008)	(0.160)	(0.013)	(0.276)	(0.027)	(0.552)
Week -4 before	0.010***	-0.03	0.019**	0.332**	-0.016	-0.774***	0.007	0.553
	(0.004)	(0.081)	(0.008)	(0.162)	(0.012)	(0.249)	(0.025)	(0.520)
Week -3 before	0.008*	0.028	0.026***	0.582***	-0.002	-0.343	0.014	0.606
	(0.004)	(0.080)	(0.008)	(0.158)	(0.012)	(0.241)	(0.024)	(0.537)
Week -2 before	0.007*	0.061	0.016**	0.301*	-0.008	-0.348	-0.011	0.939*
	(0.004)	(0.078)	(0.008)	(0.159)	(0.012)	(0.235)	(0.025)	(0.506)
Week -1 before	0.008**	0.133*	0.008	0.01	0.003	-0.465**	0.003	0.233
	(0.004)	(0.077)	(0.008)	(0.164)	(0.011)	(0.215)	(0.025)	(0.513)
Week 1 after	-0.074***	-0.972***	-0.065***	-1.079***	-0.073***	-1.512***	-0.063***	-0.638
	(0.004)	(0.076)	(0.007)	(0.157)	(0.011)	(0.229)	(0.022)	(0.509)
Week 2 after	-0.090***	-1.162***	-0.076***	-1.119***	-0.091***	-1.628***	-0.100***	-0.951*
	(0.004)	(0.077)	(0.008)	(0.161)	(0.012)	(0.247)	(0.025)	(0.554)
Week 3 after	-0.097***	-1.230***	-0.072***	-1.133***	-0.095***	-2.006***	-0.077***	-0.948*
	(0.004)	(0.078)	(0.008)	(0.169)	(0.012)	(0.244)	(0.024)	(0.554)
Week 4 after	-0.095***	-1.252***	-0.082***	-1.180***	-0.102***	-1.834***	-0.083***	-1.199**
	(0.004)	(0.082)	(0.008)	(0.166)	(0.012)	(0.233)	(0.023)	(0.547)
Week 5 after	-0.095***	-1.365***	-0.071***	-1.020***	-0.096***	-1.805***	-0.080***	-0.492
	(0.004)	(0.081)	(0.008)	(0.168)	(0.012)	(0.235)	(0.022)	(0.535)
Week 6 after	-0.097***	-1.263***	-0.085***	-1.288***	-0.093***	-1.828***	-0.073***	-1.286**
	(0.004)	(0.083)	(0.008)	(0.173)	(0.013)	(0.257)	(0.026)	(0.593)
Week 7 after	-0.096***	-1.302***	-0.081***	-1.264***	-0.092***	-1.564***	-0.079***	-0.651
	(0.004)	(0.081)	(0.008)	(0.168)	(0.013)	(0.231)	(0.026)	(0.597)
Week 8 after	-0.095***	-1.224***	-0.080***	-1.242***	-0.107***	-1.900***	-0.097***	-1.022*
	(0.004)	(0.086)	(0.008)	(0.178)	(0.012)	(0.241)	(0.024)	(0.558)
Week 9 after	-0.097***	-1.242***	-0.074***	-0.836***	-0.103***	-2.000***	-0.067***	-0.682
	(0.004)	(0.082)	(0.008)	(0.172)	(0.012)	(0.249)	(0.023)	(0.547)
Week 10 after	-0.100***	-1.351***	-0.072***	-1.083***	-0.098***	-1.838***	-0.092***	-1.232**
	(0.004)	(0.084)	(0.008)	(0.167)	(0.013)	(0.254)	(0.024)	(0.592)
Week 11 after	-0.099***	-1.257***	-0.087***	-1.181***	-0.100***	-1.591***	-0.063***	-0.841
	(0.004)	(0.083)	(0.008)	(0.170)	(0.012)	(0.235)	(0.023)	(0.619)
Week 12 after	-0.099***	-1.477***	-0.081***	-1.209***	-0.109***	-1.767***	-0.094***	-0.559
	(0.004)	(0.086)	(0.008)	(0.172)	(0.013)	(0.248)	(0.023)	(0.545)
Week 13 after	-0.098***	-1.250***	-0.080***	-1.389***	-0.110***	-1.755***	-0.068***	-0.669
	(0.004)	(0.085)	(0.008)	(0.172)	(0.013)	(0.256)	(0.024)	(0.585)
Week 14 after	-0.102***	-1.283***	-0.087***	-1.359***	-0.116***	-1.938***	-0.079***	-0.814
	(0.004)	(0.085)	(0.009)	(0.172)	(0.012)	(0.256)	(0.024)	(0.565)
Week 15 after	-0.099***	-1.349***	-0.089***	-1.488***	-0.111***	-1.956***	-0.103***	-0.837
	(0.004)	(0.085)	(0.008)	(0.175)	(0.013)	(0.253)	(0.022)	(0.552)
Week 16 after	-0.109***	-1.447***	-0.081***	-1.300***	-0.110***	-1.972***	-0.114***	-1.746***
	(0.004)	(0.084)	(0.009)	(0.187)	(0.014)	(0.272)	(0.024)	(0.586)
Week 17 after	-0.101***	-1.345***	-0.084***	-1.306***	-0.108***	-1.922***	-0.096***	-1.124*
	(0.004)	(0.087)	(0.009)	(0.185)	(0.013)	(0.251)	(0.028)	(0.634)
Week 18 after	-0.108***	-1.370***	-0.089***	-1.321***	-0.093***	-1.983***	-0.104***	-1.067*
	(0.004)	(0.086)	(0.009)	(0.188)	(0.013)	(0.265)	(0.025)	(0.558)
Week 19 after	-0.110***	-1.509***	-0.083***	-1.336***	-0.112***	-1.904***	-0.099***	-1.515***
	(0.004)	(0.086)	(0.008)	(0.180)	(0.013)	(0.272)	(0.025)	(0.548)
Week 20 after	-0.108***	-1.436***	-0.074***	-1.229***	-0.102***	-1.994***	-0.104***	-1.112**
	(0.004)	(0.086)	(0.009)	(0.183)	(0.013)	(0.261)	(0.026)	(0.504)
Pre-ticket mean	0.27	44.858	0.271	45.033	0.279	45.753	0.271	45.049
Observations	626,430	626,430	157,098	157,098	65,606	65,606	19,523	19,523
No. of cars	16,407	16,407	2,566	2,566	2,107	2,107	402	402
R2	0.233	0.243	0.230	0.236	0.241	0.267	0.243	0.280

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

	(1)	(2)	(3)	(4)
	first ticket		second ticket	
	Speeding	Speed	Speeding	Speed
Month -5 before	-0.017***	-0.455***	0.006	0.011
	(0.004)	(0.083)	(0.008)	(0.158)
Month -4 before	-0.008**	-0.262***	0.009	-0.067
	(0.004)	(0.077)	(0.007)	(0.151)
Month -3 before	-0.007**	-0.289***	0.011	0.009
	(0.003)	(0.074)	(0.007)	(0.151)
Month -2 before	-0.003	-0.099	0.007	0.026
	(0.003)	(0.071)	(0.007)	(0.137)
Month -1 before	-0.001	-0.157**	0.012*	0.087
	(0.003)	(0.067)	(0.006)	(0.129)
Month 1 after	-0.098***	-1.200***	-0.086***	-1.266***
	(0.003)	(0.070)	(0.007)	(0.146)
Month 2 after	-0.099***	-1.251***	-0.092***	-1.259***
	(0.004)	(0.072)	(0.007)	(0.147)
Month 3 after	-0.100***	-1.331***	-0.083***	-1.204***
	(0.004)	(0.074)	(0.008)	(0.141)
Month 4 after	-0.105***	-1.383***	-0.085***	-1.330***
	(0.004)	(0.071)	(0.008)	(0.161)
Month 5 after	-0.107***	-1.491***	-0.085***	-1.311***
	(0.003)	(0.073)	(0.007)	(0.154)
Month 6 after	-0.103***	-1.327***	-0.094***	-1.266***
	(0.004)	(0.075)	(0.008)	(0.144)
Month 7 after	-0.110***	-1.372***	-0.091***	-1.247***
	(0.004)	(0.076)	(0.008)	(0.148)
Month 8 after	-0.111***	-1.298***	-0.105***	-1.458***
	(0.004)	(0.078)	(0.008)	(0.148)
Month 9 after	-0.112***	-1.324***	-0.107***	-1.361***
	(0.004)	(0.077)	(0.008)	(0.148)
Month 10 after	-0.114***	-1.377***	-0.110***	-1.567***
	(0.004)	(0.078)	(0.008)	(0.146)
Month 11 after	-0.116***	-1.435***	-0.102***	-1.404***
	(0.004)	(0.077)	(0.008)	(0.157)
Month 12 after	-0.123***	-1.495***	-0.110***	-1.625***
	(0.004)	(0.076)	(0.008)	(0.162)
Month 13 after	-0.126***	-1.628***	-0.109***	-1.652***
	(0.004)	(0.080)	(0.007)	(0.153)
Month 14 after	-0.128***	-1.726***	-0.111***	-1.815***
	(0.004)	(0.079)	(0.008)	(0.154)
Month 15 after	-0.135***	-1.707***	-0.115***	-1.643***
	(0.004)	(0.080)	(0.007)	(0.150)
Month 16 after	-0.132***	-1.684***	-0.117***	-1.771***
	(0.004)	(0.078)	(0.008)	(0.155)
Month 17 after	-0.134***	-1.723***	-0.115***	-1.583***
	(0.004)	(0.081)	(0.008)	(0.155)
Month 18 after	-0.133***	-1.677***	-0.113***	-1.535***
	(0.004)	(0.081)	(0.008)	(0.158)
Month 19 after	-0.136***	-1.713***	-0.117***	-1.597***
	(0.004)	(0.082)	(0.008)	(0.162)
Month 20 after	-0.137***	-1.795***	-0.129***	-1.829***
	(0.004)	(0.087)	(0.008)	(0.160)
Month 21 after	-0.137***	-1.879***	-0.130***	-1.904***
	(0.004)	(0.092)	(0.008)	(0.168)
Month 22 after	-0.135***	-1.778***	-0.119***	-1.897***
	(0.004)	(0.095)	(0.009)	(0.182)
Month 23 after	-0.135***	-1.751***	-0.131***	-1.896***
	(0.004)	(0.092)	(0.009)	(0.184)
Month 24 after	-0.139***	-1.869***	-0.127***	-1.753***
	(0.004)	(0.096)	(0.009)	(0.189)
Pre-ticket mean	0.257	44.353	0.275	44.786
Observations	991,333	991,333	258,540	258,540
No. of cars	4291	4291	891	891
R2	0.188	0.204	0.190	0.207

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)
	Car Owner				*Payment of Ticket*			
	private	corporation	private	corporation	unpaid	paid	unpaid	paid
	Speeding		Speed		Speeding		Speed	
Week -11 before	0.002	0.002	0.12	-0.05	0.026**	-0.002	0.31	-0.01
	(0.006)	(0.006)	(0.12)	(0.13)	(0.011)	(0.005)	(0.22)	(0.09)
Week -10 before	0.006	-0.005	0.12	-0.28**	0.020*	-0.003	0.16	-0.12
	(0.006)	(0.006)	(0.12)	(0.13)	(0.010)	(0.005)	(0.24)	(0.10)
Week -9 before	0.013**	-0.006	0.07	-0.16	0.027**	-0.001	0.05	-0.06
	(0.006)	(0.006)	(0.12)	(0.13)	(0.011)	(0.005)	(0.23)	(0.09)
Week -8 before	0.007	-0.007	-0.03	-0.18	0.006	-0.001	-0.03	-0.11
	(0.006)	(0.006)	(0.12)	(0.13)	(0.010)	(0.005)	(0.22)	(0.09)
Week -7 before	0.004	-0.013**	0.05	-0.37***	0.016	-0.009*	0.38*	-0.26***
	(0.006)	(0.006)	(0.12)	(0.12)	(0.011)	(0.004)	(0.21)	(0.09)
Week -6 before	0.010*	-0.005	-0.06	-0.1	0.013	0.001	0.27	-0.14
	(0.006)	(0.006)	(0.11)	(0.12)	(0.010)	(0.004)	(0.21)	(0.09)
Week -5 before	0.009	-0.007	-0.04	-0.18	0.017*	-0.002	0.1	-0.14
	(0.006)	(0.006)	(0.11)	(0.12)	(0.010)	(0.004)	(0.21)	(0.09)
Week -4 before	0.013**	0.008	0.09	-0.16	0.030***	0.006	0.42**	-0.12
	(0.006)	(0.006)	(0.11)	(0.12)	(0.010)	(0.004)	(0.21)	(0.09)
Week -3 before	0.009	0.006	0.05	-0.01	0.021**	0.005	0.43**	-0.05
	(0.006)	(0.006)	(0.11)	(0.11)	(0.010)	(0.004)	(0.20)	(0.09)
Week -2 before	0.009*	0.005	0.12	-0.01	0.022**	0.004	0.27	0.02
	(0.005)	(0.005)	(0.11)	(0.11)	(0.009)	(0.004)	(0.20)	(0.08)
Week -1 before	0.014***	0.001	0.26**	0	0.015	0.006	0.26	0.11
	(0.005)	(0.005)	(0.10)	(0.11)	(0.009)	(0.004)	(0.20)	(0.08)
Week 1 after	-0.086***	-0.062***	-1.26***	-0.70***	-0.030***	-0.082***	-0.33*	-1.10***
	(0.005)	(0.005)	(0.10)	(0.11)	(0.009)	(0.004)	(0.19)	(0.08)
Week 2 after	-0.099***	-0.082***	-1.35***	-0.99***	-0.040***	-0.100***	-0.70***	-1.25***
	(0.005)	(0.005)	(0.10)	(0.11)	(0.009)	(0.004)	(0.20)	(0.08)
Week 3 after	-0.109***	-0.086***	-1.47***	-1.00***	-0.049***	-0.107***	-0.42**	-1.39***
	(0.005)	(0.006)	(0.11)	(0.11)	(0.010)	(0.004)	(0.20)	(0.08)
Week 4 after	-0.105***	-0.085***	-1.35***	-1.16***	-0.054***	-0.103***	-0.66***	-1.37***
	(0.005)	(0.006)	(0.11)	(0.12)	(0.010)	(0.004)	(0.21)	(0.09)
Week 5 after	-0.107***	-0.084***	-1.53***	-1.21***	-0.038***	-0.107***	-0.65***	-1.51***
	(0.005)	(0.006)	(0.11)	(0.12)	(0.010)	(0.004)	(0.20)	(0.09)
Week 6 after	-0.107***	-0.088***	-1.42***	-1.11***	-0.048***	-0.107***	-0.60***	-1.39***
	(0.006)	(0.006)	(0.11)	(0.12)	(0.010)	(0.004)	(0.22)	(0.09)
Week 7 after	-0.106***	-0.087***	-1.43***	-1.19***	-0.051***	-0.106***	-0.96***	-1.37***
	(0.006)	(0.006)	(0.11)	(0.12)	(0.010)	(0.004)	(0.22)	(0.09)
Week 8 after	-0.102***	-0.089***	-1.37***	-1.09***	-0.048***	-0.105***	-0.53**	-1.36***
	(0.006)	(0.006)	(0.11)	(0.13)	(0.011)	(0.004)	(0.22)	(0.09)
Week 9 after	-0.101***	-0.093***	-1.45***	-1.04***	-0.055***	-0.106***	-0.74***	-1.34***
	(0.006)	(0.006)	(0.11)	(0.12)	(0.011)	(0.004)	(0.22)	(0.09)
Week 10 after	-0.107***	-0.094***	-1.47***	-1.24***	-0.063***	-0.107***	-0.82***	-1.46***
	(0.006)	(0.006)	(0.11)	(0.12)	(0.010)	(0.004)	(0.22)	(0.09)
Week 11 after	-0.110***	-0.088***	-1.44***	-1.09***	-0.055***	-0.107***	-0.81***	-1.34***
	(0.006)	(0.006)	(0.11)	(0.12)	(0.011)	(0.004)	(0.23)	(0.09)
Week 12 after	-0.108***	-0.090***	-1.60***	-1.37***	-0.059***	-0.107***	-1.19***	-1.54***
	(0.006)	(0.006)	(0.12)	(0.12)	(0.011)	(0.004)	(0.24)	(0.09)
Week 13 after	-0.105***	-0.093***	-1.41***	-1.09***	-0.068***	-0.104***	-0.80***	-1.34***
	(0.006)	(0.006)	(0.12)	(0.12)	(0.011)	(0.004)	(0.23)	(0.09)
Week 14 after	-0.108***	-0.097***	-1.47***	-1.11***	-0.065***	-0.109***	-0.91***	-1.36***
	(0.006)	(0.006)	(0.12)	(0.12)	(0.011)	(0.004)	(0.24)	(0.09)
Week 15 after	-0.107***	-0.092***	-1.57***	-1.15***	-0.070***	-0.105***	-0.94***	-1.43***
	(0.006)	(0.006)	(0.11)	(0.12)	(0.011)	(0.004)	(0.22)	(0.09)
Week 16 after	-0.113***	-0.106***	-1.62***	-1.28***	-0.089***	-0.114***	-0.97***	-1.54***
	(0.006)	(0.006)	(0.11)	(0.12)	(0.011)	(0.005)	(0.22)	(0.09)
Week 17 after	-0.106***	-0.096***	-1.38***	-1.32***	-0.090***	-0.103***	-1.12***	-1.39***
	(0.006)	(0.006)	(0.12)	(0.13)	(0.011)	(0.005)	(0.22)	(0.09)
Week 18 after	-0.114***	-0.103***	-1.53***	-1.23***	-0.080***	-0.115***	-0.94***	-1.46***
	(0.006)	(0.006)	(0.12)	(0.12)	(0.011)	(0.004)	(0.23)	(0.09)
Week 19 after	-0.115***	-0.105***	-1.62***	-1.41***	-0.094***	-0.114***	-1.29***	-1.55***
	(0.006)	(0.006)	(0.12)	(0.12)	(0.011)	(0.004)	(0.23)	(0.09)
Week 20 after	-0.109***	-0.107***	-1.49***	-1.40***	-0.078***	-0.114***	-1.00***	-1.52***
	(0.006)	(0.006)	(0.12)	(0.13)	(0.012)	(0.004)	(0.23)	(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)
	by driving frequency				*by number plate region*					
	Infrequent	Frequent	Infrequent	Frequent	Local region	Prague	Other regions	Local region	Prague	Other regions
	Speeding		Speed		Speeding			Speed		
Week -11 before	0.006	0.002	0.28	0.02	-0.006	0.011*	0.001	-0.1	0.19	0.03
	(0.013)	(0.005)	(0.24)	(0.09)	(0.006)	(0.006)	(0.013)	(0.13)	(0.13)	(0.26)
Week -10 before	-0.001	0.001	0.13	-0.09	-0.001	0.002	0.001	-0.04	-0.11	-0.10
	(0.013)	(0.004)	(0.24)	(0.10)	(0.006)	(0.006)	(0.012)	(0.13)	(0.14)	(0.26)
Week -9 before	0.013	0.004	-0.16	-0.02	0.007	-0.002	0.015	0.07	-0.18	0.08
	(0.013)	(0.004)	(0.24)	(0.09)	(0.006)	(0.006)	(0.013)	(0.13)	(0.13)	(0.26)
Week -8 before	0.001	0.001	0.05	-0.11	0	0.003	-0.008	-0.1	-0.1	-0.13
	(0.013)	(0.004)	(0.24)	(0.09)	(0.006)	(0.006)	(0.012)	(0.13)	(0.13)	(0.26)
Week -7 before	0.009	-0.005	-0.26	-0.14	-0.006	-0.003	-0.004	-0.11	-0.18	-0.20
	(0.012)	(0.004)	(0.24)	(0.09)	(0.006)	(0.006)	(0.011)	(0.13)	(0.13)	(0.24)
Week -6 before	0.012	0.002	0.02	-0.08	-0.003	0.005	0.015	-0.14	-0.04	0.00
	(0.012)	(0.004)	(0.22)	(0.09)	(0.006)	(0.006)	(0.012)	(0.13)	(0.12)	(0.24)
Week -5 before	-0.01	0.003	-0.22	-0.08	0.005	-0.003	0.004	-0.17	-0.01	-0.19
	(0.012)	(0.004)	(0.22)	(0.09)	(0.006)	(0.006)	(0.012)	(0.12)	(0.12)	(0.24)
Week -4 before	0.025**	0.009**	0.02	-0.04	0.007	0.014**	0.01	0.01	0.01	-0.27
	(0.011)	(0.004)	(0.22)	(0.09)	(0.006)	(0.006)	(0.011)	(0.12)	(0.12)	(0.24)
Week -3 before	0.008	0.008*	0.21	0	0.004	0.009	0.015	-0.07	0.11	0.05
	(0.011)	(0.004)	(0.21)	(0.09)	(0.006)	(0.006)	(0.011)	(0.12)	(0.12)	(0.23)
Week -2 before	0.018	0.006	0.28	0.03	0.005	0.007	0.017	-0.06	0.18	0.07
	(0.011)	(0.004)	(0.21)	(0.08)	(0.006)	(0.006)	(0.011)	(0.12)	(0.11)	(0.22)
Week -1 before	0.016	0.007*	0.23	0.11	0.006	0.004	0.025**	0.12	0.16	0.06
	(0.010)	(0.004)	(0.20)	(0.08)	(0.006)	(0.006)	(0.010)	(0.12)	(0.12)	(0.21)
Week 1 after	-0.085***	-0.072***	-1.20***	-0.93***	-0.085***	-0.063***	-0.071***	-1.10***	-0.82***	-1.10***
	(0.010)	(0.004)	(0.20)	(0.08)	(0.006)	(0.006)	(0.011)	(0.11)	(0.12)	(0.22)
Week 2 after	-0.112***	-0.087***	-1.47***	-1.11***	-0.095***	-0.082***	-0.105***	-1.28***	-1.07***	-1.12***
	(0.010)	(0.004)	(0.20)	(0.08)	(0.006)	(0.006)	(0.011)	(0.11)	(0.12)	(0.20)
Week 3 after	-0.123***	-0.093***	-1.48***	-1.19***	-0.106***	-0.083***	-0.113***	-1.42***	-1.00***	-1.42***
	(0.010)	(0.004)	(0.20)	(0.08)	(0.006)	(0.006)	(0.011)	(0.11)	(0.12)	(0.21)
Week 4 after	-0.097***	-0.094***	-1.43***	-1.22***	-0.100***	-0.087***	-0.104***	-1.38***	-1.07***	-1.47***
	(0.010)	(0.004)	(0.20)	(0.09)	(0.006)	(0.006)	(0.011)	(0.12)	(0.12)	(0.22)
Week 5 after	-0.110***	-0.093***	-1.64***	-1.32***	-0.106***	-0.081***	-0.109***	-1.51***	-1.15***	-1.70***
	(0.010)	(0.004)	(0.21)	(0.09)	(0.006)	(0.006)	(0.011)	(0.12)	(0.12)	(0.22)
Week 6 after	-0.110***	-0.095***	-1.36***	-1.25***	-0.106***	-0.084***	-0.108***	-1.50***	-1.01***	-1.29***
	(0.011)	(0.004)	(0.21)	(0.09)	(0.006)	(0.006)	(0.011)	(0.12)	(0.12)	(0.22)
Week 7 after	-0.108***	-0.095***	-1.57***	-1.26***	-0.106***	-0.086***	-0.098***	-1.39***	-1.14***	-1.54***
	(0.011)	(0.004)	(0.21)	(0.09)	(0.006)	(0.006)	(0.012)	(0.12)	(0.12)	(0.23)
Week 8 after	-0.110***	-0.093***	-1.37***	-1.20***	-0.102***	-0.083***	-0.120***	-1.34***	-0.96***	-1.77***
	(0.011)	(0.004)	(0.20)	(0.09)	(0.006)	(0.006)	(0.012)	(0.13)	(0.13)	(0.24)
Week 9 after	-0.111***	-0.095***	-1.40***	-1.22***	-0.105***	-0.082***	-0.128***	-1.48***	-0.86***	-1.76***
	(0.011)	(0.004)	(0.21)	(0.09)	(0.006)	(0.006)	(0.012)	(0.12)	(0.12)	(0.23)
Week 10 after	-0.123***	-0.097***	-1.75***	-1.30***	-0.104***	-0.091***	-0.122***	-1.43***	-1.17***	-1.77***
	(0.011)	(0.004)	(0.22)	(0.09)	(0.006)	(0.006)	(0.013)	(0.12)	(0.12)	(0.26)
Week 11 after	-0.109***	-0.097***	-1.32***	-1.25***	-0.107***	-0.086***	-0.115***	-1.35***	-1.12***	-1.47***
	(0.011)	(0.004)	(0.21)	(0.09)	(0.006)	(0.006)	(0.012)	(0.12)	(0.13)	(0.24)
Week 12 after	-0.112***	-0.097***	-1.55***	-1.47***	-0.108***	-0.084***	-0.122***	-1.55***	-1.38***	-1.57***
	(0.011)	(0.004)	(0.22)	(0.09)	(0.006)	(0.006)	(0.012)	(0.13)	(0.13)	(0.25)
Week 13 after	-0.127***	-0.094***	-1.75***	-1.17***	-0.104***	-0.086***	-0.122***	-1.32***	-1.15***	-1.31***
	(0.011)	(0.005)	(0.21)	(0.09)	(0.006)	(0.006)	(0.012)	(0.13)	(0.13)	(0.24)
Week 14 after	-0.116***	-0.100***	-1.62***	-1.23***	-0.113***	-0.088***	-0.113***	-1.43***	-1.06***	-1.49***
	(0.011)	(0.004)	(0.21)	(0.09)	(0.006)	(0.006)	(0.012)	(0.13)	(0.13)	(0.23)
Week 15 after	-0.113***	-0.097***	-1.60***	-1.31***	-0.109***	-0.088***	-0.106***	-1.45***	-1.20***	-1.53***
	(0.011)	(0.005)	(0.22)	(0.09)	(0.006)	(0.006)	(0.012)	(0.12)	(0.13)	(0.22)
Week 16 after	-0.121***	-0.107***	-1.65***	-1.41***	-0.113***	-0.102***	-0.120***	-1.61***	-1.22***	-1.67***
	(0.011)	(0.005)	(0.21)	(0.09)	(0.006)	(0.006)	(0.012)	(0.12)	(0.13)	(0.24)
Week 17 after	-0.112***	-0.099***	-1.74***	-1.28***	-0.105***	-0.092***	-0.116***	-1.41***	-1.22***	-1.60***
	(0.011)	(0.005)	(0.22)	(0.09)	(0.006)	(0.007)	(0.012)	(0.13)	(0.13)	(0.24)
Week 18 after	-0.124***	-0.106***	-1.73***	-1.31***	-0.113***	-0.098***	-0.128***	-1.36***	-1.31***	-1.73***
	(0.011)	(0.004)	(0.22)	(0.09)	(0.006)	(0.006)	(0.012)	(0.13)	(0.13)	(0.26)
Week 19 after	-0.137***	-0.106***	-1.85***	-1.46***	-0.113***	-0.103***	-0.128***	-1.50***	-1.39***	-2.06***
	(0.011)	(0.004)	(0.22)	(0.09)	(0.006)	(0.006)	(0.012)	(0.13)	(0.13)	(0.24)
Week 20 after	-0.116***	-0.106***	-1.77***	-1.39***	-0.113***	-0.095***	-0.133***	-1.53***	-1.20***	-1.96***
	(0.011)	(0.005)	(0.22)	(0.09)	(0.006)	(0.007)	(0.012)	(0.13)	(0.13)	(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: by drive-through order

	(1)	(2)	(3)	(4)
	First Ticket		Second Ticket	
	Speeding	Speed	Speeding	Speed
Drive -61 to -65	-0.018***	-0.26***	-0.001	-0.05
	(0.003)	(0.07)	(0.007)	(0.13)
Drive -56 to 60	-0.020***	-0.43***	-0.008	-0.30**
	(0.003)	(0.07)	(0.007)	(0.13)
Drive -50 to -55	-0.017***	-0.36***	0.003	-0.18
	(0.003)	(0.07)	(0.006)	(0.13)
Drive -46 to -50	-0.020***	-0.40***	0.005	-0.01
	(0.003)	(0.07)	(0.006)	(0.13)
Drive -41 to -45	-0.020***	-0.34***	0.002	-0.25**
	(0.003)	(0.06)	(0.006)	(0.13)
Drive -36 to -40	-0.015***	-0.30***	0.007	0.20*
	(0.003)	(0.06)	(0.006)	(0.12)
Drive -31 to -35	-0.008***	-0.27***	-0.004	-0.11
	(0.003)	(0.06)	(0.006)	(0.12)
Drive -26 to -30	-0.011***	-0.21***	0.006	0.08
	(0.003)	(0.06)	(0.006)	(0.12)
Drive -21 to -25	-0.010***	-0.23***	0.005	-0.02
	(0.003)	(0.05)	(0.006)	(0.11)
Drive -16 to -20	-0.007***	-0.13***	0.003	0.15
	(0.003)	(0.05)	(0.005)	(0.11)
Drive -11 to -15	-0.001	-0.04	0.003	0.04
	(0.002)	(0.05)	(0.005)	(0.11)
Drive -6 to -10	-0.000	0.02	0.006	0.14
	(0.002)	(0.04)	(0.005)	(0.10)
Drive 1 to 5	-0.093***	-1.16***	-0.095***	-1.46***
	(0.002)	(0.04)	(0.005)	(0.10)
Drive 6 to 10	-0.105***	-1.31***	-0.107***	-1.72***
	(0.002)	(0.04)	(0.005)	(0.11)
Drive 11 to 15	-0.110***	-1.34***	-0.110***	-1.68***
	(0.002)	(0.05)	(0.006)	(0.11)
Drive 16 to 20	-0.115***	-1.47***	-0.116***	-1.79***
	(0.002)	(0.05)	(0.006)	(0.11)
Drive 21 to 25	-0.117***	-1.49***	-0.117***	-1.82***
	(0.003)	(0.05)	(0.006)	(0.11)
Drive 26 to 30	-0.120***	-1.55***	-0.107***	-1.66***
	(0.003)	(0.05)	(0.006)	(0.12)
Drive 31 to 35	-0.120***	-1.52***	-0.110***	-1.83***
	(0.003)	(0.06)	(0.006)	(0.12)
Drive 36 to 40	-0.123***	-1.56***	-0.117***	-1.77***
	(0.003)	(0.06)	(0.006)	(0.12)
Drive 41 to 45	-0.123***	-1.55***	-0.118***	-1.87***
	(0.003)	(0.06)	(0.006)	(0.12)
Drive 46 to 50	-0.123***	-1.58***	-0.105***	-1.66***
	(0.003)	(0.06)	(0.006)	(0.13)
Drive 51 to 55	-0.121***	-1.55***	-0.111***	-1.65***
	(0.003)	(0.06)	(0.006)	(0.13)
Drive 56 to 60	-0.128***	-1.61***	-0.111***	-1.75***
	(0.003)	(0.06)	(0.007)	(0.13)
Drive 61 to 65	-0.126***	-1.57***	-0.110***	-1.78***
	(0.003)	(0.06)	(0.007)	(0.14)
Drive 66 to 70	-0.126***	-1.59***	-0.109***	-1.72***
	(0.003)	(0.06)	(0.007)	(0.14)
Drive 71 to 75	-0.125***	-1.52***	-0.117***	-1.97***
	(0.003)	(0.06)	(0.007)	(0.14)
Drive 76 to 80	-0.127***	-1.68***	-0.116***	-1.64***
	(0.003)	(0.06)	(0.007)	(0.14)
Drive 81 to 85	-0.131***	-1.60***	-0.117***	-1.78***
	(0.003)	(0.07)	(0.007)	(0.14)
Drive 86 to 90	-0.129***	-1.61***	-0.121***	-1.88***
	(0.003)	(0.07)	(0.007)	(0.15)
Drive 91 to 95	-0.132***	-1.64***	-0.117***	-1.87***
	(0.003)	(0.07)	(0.007)	(0.15)
Drive 96 to 100	-0.129***	-1.69***	-0.116***	-1.67***
	(0.003)	(0.07)	(0.007)	(0.15)
Pre-ticket mean	0.299	-3.665	0.326	-2.898
Observations	1,171,931	1,171,931	260,513	260,513
No. of cars	16,414	16,414	2,566	2,566
R2	0.213	0.22	0.221	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.

Table A.16: Event analysis: same versus other speed camera zone

	(1)	(2)	(3)	(4)
	Same zone	Other zones	Same zone	Other zones
	Speeding		Speed	
Week -11 before	0.007	0.004	-0.093	0.201*
	(0.007)	(0.005)	(0.143)	(0.109)
Week -10 before	0.004	0.005	-0.269*	0.155
	(0.007)	(0.005)	(0.151)	(0.105)
Week -9 before	0.009	0.006	-0.010	0.039
	(0.007)	(0.005)	(0.141)	(0.106)
Week -8 before	0.002	0.005	-0.256*	0.100
	(0.007)	(0.005)	(0.144)	(0.104)
Week -7 before	0.003	-0.003	-0.269*	0.027
	(0.007)	(0.005)	(0.141)	(0.104)
Week -6 before	0.004	0.005	-0.123	-0.008
	(0.007)	(0.005)	(0.136)	(0.105)
Week -5 before	-0.002	0.005	-0.053	-0.093
	(0.007)	(0.005)	(0.131)	(0.101)
Week -4 before	0.010	0.011**	-0.002	-0.015
	(0.007)	(0.005)	(0.129)	(0.102)
Week -3 before	0.009	0.008*	0.034	0.054
	(0.007)	(0.005)	(0.126)	(0.100)
Week -2 before	0.009	0.008*	0.026	0.124
	(0.006)	(0.005)	(0.126)	(0.095)
Week -1 before	0.009	0.008*	0.104	0.173*
	(0.006)	(0.005)	(0.123)	(0.095)
Week 1 after	-0.102***	-0.053***	-1.394***	-0.640***
	(0.006)	(0.004)	(0.121)	(0.096)
Week 2 after	-0.125***	-0.064***	-1.620***	-0.803***
	(0.006)	(0.004)	(0.124)	(0.095)
Week 3 after	-0.139***	-0.066***	-1.714***	-0.840***
	(0.007)	(0.004)	(0.124)	(0.096)
Week 4 after	-0.137***	-0.064***	-1.765***	-0.868***
	(0.007)	(0.005)	(0.131)	(0.100)
Week 5 after	-0.136***	-0.069***	-2.003***	-0.889***
	(0.007)	(0.005)	(0.129)	(0.098)
Week 6 after	-0.136***	-0.070***	-1.780***	-0.841***
	(0.007)	(0.005)	(0.133)	(0.099)
Week 7 after	-0.145***	-0.063***	-1.989***	-0.796***
	(0.007)	(0.005)	(0.133)	(0.100)
Week 8 after	-0.141***	-0.066***	-1.885***	-0.744***
	(0.007)	(0.005)	(0.140)	(0.103)
Week 9 after	-0.141***	-0.068***	-1.835***	-0.812***
	(0.007)	(0.005)	(0.133)	(0.100)
Week 10 after	-0.140***	-0.073***	-1.924***	-0.932***
	(0.007)	(0.005)	(0.135)	(0.101)
Week 11 after	-0.142***	-0.068***	-1.867***	-0.809***
	(0.007)	(0.005)	(0.132)	(0.103)
Week 12 after	-0.150***	-0.062***	-2.028***	-1.045***
	(0.007)	(0.005)	(0.139)	(0.106)
Week 13 after	-0.141***	-0.070***	-1.858***	-0.840***
	(0.007)	(0.005)	(0.137)	(0.105)
Week 14 after	-0.146***	-0.071***	-1.819***	-0.890***
	(0.007)	(0.005)	(0.135)	(0.104)
Week 15 after	-0.143***	-0.068***	-1.999***	-0.855***
	(0.007)	(0.005)	(0.135)	(0.102)
Week 16 after	-0.156***	-0.075***	-2.120***	-0.892***
	(0.007)	(0.005)	(0.135)	(0.104)
Week 17 after	-0.149***	-0.065***	-2.085***	-0.769***
	(0.007)	(0.005)	(0.141)	(0.103)
Week 18 after	-0.155***	-0.075***	-2.032***	-0.880***
	(0.007)	(0.005)	(0.134)	(0.107)
Week 19 after	-0.159***	-0.078***	-2.239***	-0.996***
	(0.007)	(0.005)	(0.137)	(0.106)
Week 20 after	-0.152***	-0.077***	-2.138***	-0.902***
	(0.007)	(0.005)	(0.144)	(0.105)
Pre-ticket mean	0.362	0.199	46.951	43.244
Observations	262,282	361,352	262,282	361,352
No. of cars	13,769	14,104	13,769	14,104
R2	0.273	0.233	0.293	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.

B Complementary Online Appendix

B.1 Mapping experiences in expectations

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

$$\begin{aligned} p^t(s) &= P\left(\{s_{t-1}, \mathrm{T}^t(s_{t-1})\}, \{s_{t-2}, \mathrm{T}^t(s_{t-2})\}, \ldots, \{s_0, \mathrm{T}^t(s_0)\},\right. \\ &\quad \left. P\left(\{s_{t-2}, \mathrm{T}^{t-1}(s_{t-2})\}, \ldots, \{s_0, \mathrm{T}^{t-1}(s_0)\},\ p^{t-2}(s)\right)\right) = \ldots \end{aligned} \tag{B.1}$$

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

$$\begin{aligned} &= \Pi_t\left(\{s_{t-1}, \mathrm{T}^t(s_{t-1})\}, \{s_{t-2}, \mathrm{T}^t(s_{t-2}), \mathrm{T}^{t-1}(s_{t-2})\}, \ldots,\right. \\ &\quad \left.\{s_0, \mathrm{T}^t(s_0), \mathrm{T}^{t-1}(s_0), \ldots, \mathrm{T}^1(s_0)\},\ p^0(s)\right) \end{aligned} \tag{B.2}$$

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

$$p^t(s) = \Pi_t\left(\left(\{s_{t-1}, \vec{\mathrm{T}}(t, s_{t-1})\}\right), \ldots, \left(\{s_1, \vec{\mathrm{T}}(t, s_1)\}\right), \left(\{s_0, \vec{\mathrm{T}}(t, s_0)\}\right), p^0(s)\right), \tag{B.3}$$

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 define 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), starts the first time a car is observed. For both cars in this example we observe the highest speed during the assignment period at the same day (a bit 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 is 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 would reach 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

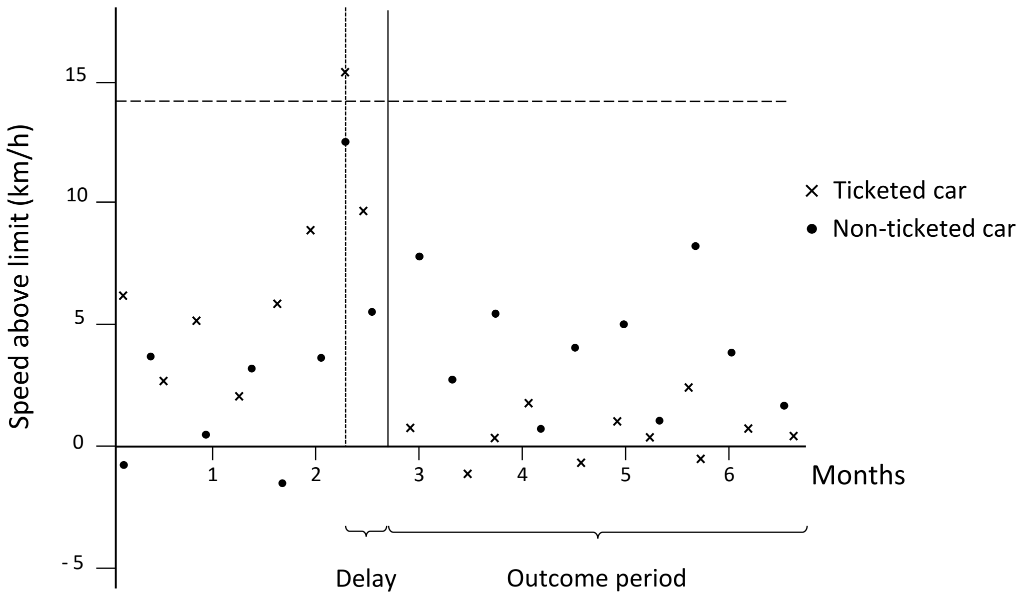

Notes: The figure illustrates our approach to define assignment and outcome periods for the driving patterns of two cars.

The Figure from 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 holds again symmetrically for cars with an assignments speed below or above any of the two RDD cutoffs. Hence, this properties 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 assignment window and latter months as outcome periods. This static approach produced very similar results, but explored a much smaller part of the sample.

B.3 Mean reversion is event analysis

To illustrate the mean reversion issue in the raw data (which is captured in Figure 9), we introduce a simple framework of speed choices. Speed s_{it} at a drive-through at time t by car i is given by

$$s_{it} = \lambda_i + \beta T_{it} + \gamma X_{it} + \varepsilon_{it}, \tag{B.4}$$

where λ_i is a car fixed effect, T_{it} is a treatment dummy (indicating that the driver has received the 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 later on. Hence, the ticket dummy is positive only if the car was driving above the enforcement cutoff k^1 at some point t'. This can be written as

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

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

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

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 in one to 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, one would overestimate the effect from the speeding tickets.