

Forgetting and Heterogeneity in Task Delay: Evidence from New York City Parking-Ticket Recipients*

Ori Heffetz, Ted O’Donoghue, and Henry S. Schneider

December 26, 2016

Abstract

We study response behavior of New York City parking-ticket recipients by analyzing administrative data on 6.6 million tickets issued to 2 million individuals over two years. Exploiting variation (from a policy change and a field experiment) in letters sent to recipients, we find that forgetting plays a major role in delay—letters seem to act mostly as reminders, with their content mattering little. Moreover, by studying an individual’s behavior across multiple tickets, we find significant heterogeneity in underlying types, with different types reacting differently to deadlines and reminders. Failure to account for this heterogeneity yields biased—and sometimes incorrect—conclusions.

*Heffetz: S.C. Johnson Graduate School of Management, Cornell University, 324 Sage Hall, Ithaca, NY 14853, and Department of Economics, Hebrew University of Jerusalem (e-mail: oh33@cornell.edu); O’Donoghue: Department of Economics, Cornell University, 482 Uris Hall, Ithaca, NY 14853 (e-mail: edo1@cornell.edu); Schneider: Stephen J.R. Smith School of Business, Queen’s University, Goodes Hall, Kingston, ON K7L 3N6 (e-mail: henry.schneider@queensu.ca). We thank Samuel Bufter, David Frankel, Christopher Quinn, and Andrew Salkin at New York City Department of Finance, and Janet Capuano at CGI Group for access to the data and many helpful conversations. We have benefitted from valuable comments from Amr Farahat and seminar participants at Ben Gurion, Cornell, Harvard, Hebrew U, IDC Herzliya, UCLA Anderson, UCSD Rady, USC, Wharton, Williams, SITE 2014, and BEAM 2015. We also thank Lawrence Jin, Lin Xu, and Maya Catabi for excellent research assistance. The authors have no financial or other material interests related to this research to disclose.

Online Appendices are available on the authors’ websites.

1 Introduction

You return to your car to find a parking ticket on the windshield. When, if ever, will you pay it? Life involves many such tasks that must—or at least should—be completed. People must pay bills, file taxes, make doctors’ appointments and fill prescriptions, apply for schools and jobs, and so forth. Despite the ubiquity of such tasks, there is relatively little research in economics on the determinants of when—and whether—people complete them.

In this paper, we study a classic task-completion problem: when to respond to a parking ticket. Using a detailed administrative dataset provided by the New York City Department of Finance (henceforth, DOF), we observe response behavior associated with the universe of parking tickets issued in New York City between June 2011 and August 2013. In our core dataset, we analyze response behavior for 6,646,540 tickets issued to 1,980,698 unique passenger vehicle license plates, totalling \$424 million in fines and \$85 million in late penalties.

Our paper makes three contributions. First, the large size of our dataset permits us to non-parametrically characterize aggregate response rates at the daily level over a very long period of time—for 135 days after people receive their tickets. As such, we provide a detailed look at behavior in this domain. Second, we focus on one particular determinant of task completion: forgetting. Our dataset contains exogenous variation in both the timing and the content of notification letters that are sent by DOF to ticket recipients, created by a combination of a citywide regime shift implemented in June 2012 and a citywide field experiment implemented in the summer of 2013. Using this variation, we find strong evidence of forgetting—in particular, notification letters seem to act mostly as reminders, with their content mattering little. Third, our analysis sheds light on the important role of heterogeneity in this domain. Unlike prior empirical work on task completion (especially that on forgetting), our dataset contains response behavior for the same individuals (license plates) across multiple tasks (tickets). Using this feature, we find clear indication of persistent types. Moreover, we find that different types react differently to deadlines and reminders. Incidence analysis of this kind is often essential for formulating effective policy interventions.

In Section 2, we develop a stylized theoretical model of task completion, designed to roughly match our field context and to help organize our interpretation of behavior. We

begin with a search model in the spirit of McCall (1970) that captures the intuition that people wait for a convenient (low-cost) time to pay their tickets. We then describe the implications of forgetting and heterogeneity in that context.

In Section 3, we describe our data and setting in more detail. After receiving a ticket, the ticket recipient faces a series of deadlines and associated late penalties. If she does not respond by a first deadline (day 30)—by either paying or contesting (i.e., challenging) the ticket—she is assessed a \$10 late penalty. If she still does not respond by a second deadline (the first Monday after day 61), she is assessed an additional \$20 penalty. If she still does not respond by a third deadline (the first Friday after day 100), she is assessed yet an additional \$30 penalty. The (day 0) windshield ticket clearly states the first deadline and penalty. At various later times, the ticket recipient receives notification letters from DOF to keep her informed of her current situation and to specify updated deadlines and penalties.

While these deadlines and late penalties did not change throughout the timeframe of our data, the timing and content of notification letters did. In particular, in June 2012, DOF implemented what we label as a shift from the *OLD regime* to the *NEW regime*. The key change was in the timing of the first notification letter: under the OLD regime, it was received at roughly day 40, whereas under the NEW regime it was received at roughly day 20. Our initial test of forgetting exploits this change.

In Section 4, we analyze response behavior in the OLD and NEW regimes. Figure 1 depicts daily hazard rates of recipients’ first responses (i.e., $\frac{\# \text{ first response on day } d}{\# \text{ no response before day } d}$) as well as cumulative response rates (i.e., $\frac{\# \text{ first response on or before day } d}{\text{total } \# \text{ of tickets issued}}$). The horizontal axis indicates the number of days since the ticket was issued, and the three deadlines are highlighted by the vertical shaded bands.¹ Figure 1 suggests that the change in timing of the first letter had a striking impact that is consistent with forgetting. Relative to the OLD regime, where no communications are received prior to the first deadline, under the NEW regime there is a dramatic increase in hazard rates following day 20. In addition, relative to the NEW regime, where no communications are received between the first and second deadlines, under the OLD regime there is a dramatic increase in hazard rates following day 40.

¹After the third deadline, DOF enters a default judgment in court against the plate owner, after which more serious actions might be taken. Figure 1 reveals that this action has a strong impact on hazard rates. While we focus on behavior prior to the third deadline, we discuss this later behavior where relevant.

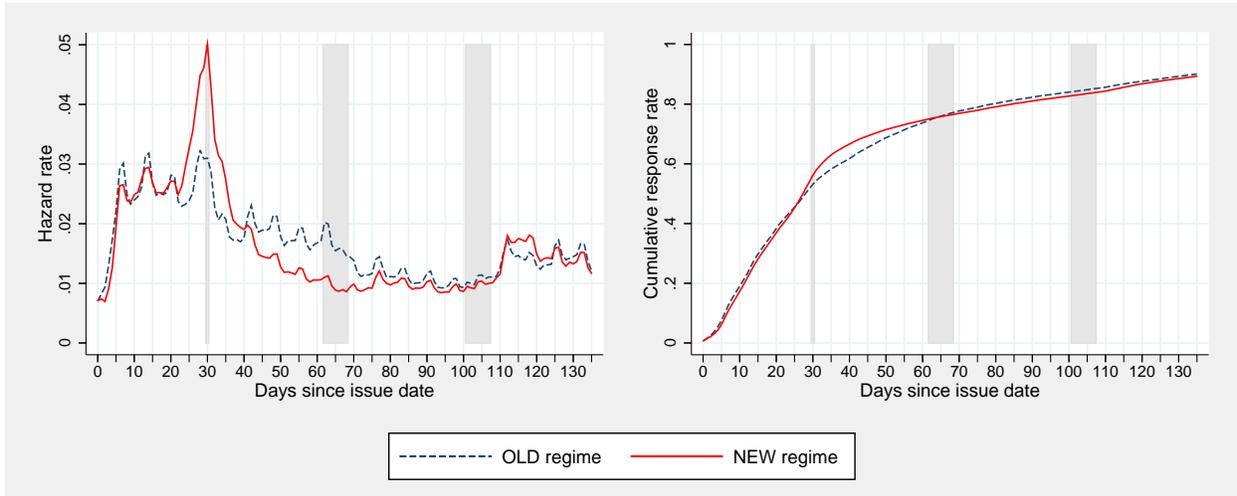


Figure 1: **Hazard Rates and Cumulative Response Rates in OLD vs. NEW Regimes.** Note: All tickets have a first deadline at day 30, second deadline at days 62-68, and third deadline at days 101-107, indicated by the shaded areas (the latter two deadlines are a range because they depend on ticket-issuance day of the week). First notification letter is received around day 40 (OLD) vs. day 20 (NEW). Based on 3,355,094 (OLD) and 3,020,357 (NEW) observations; see details in Section 3.

While the pattern in Figure 1 is consistent with forgetting, the specific way in which DOF implemented the regime change resulted in issues of interpretation. To rule out alternative explanations, we worked with DOF on a field experiment that was implemented over five weeks, for tickets issued July 13, 2013 through August 16, 2013. The recipient of each ticket was assigned to receive one of four versions of the first letter (at roughly day 20), and to receive or not receive an additional letter at roughly day 48 (a 4×2 experimental design). The four versions of the first letter vary in their informational content and presentation. The results seem to confirm that the letters served primarily as reminders.

Figure 1 exhibits other important patterns. First, under both regimes, there is a dramatic decline in hazard rates between the first and second deadlines. While this decline could be due in part to forgetting, a similar decline in hazard rates over time is typically observed in job search data, and has been attributed to heterogeneity in types (Salant (1977)). Second, while there is a noticeable spike in hazard rates at the first deadline under both regimes (and especially under the NEW regime), there are at most small spikes in hazard rates at the second and third deadlines. This apparent limited response to later deadlines is also

consistent with heterogeneity in types: if only some types pay attention to deadlines, those types disproportionately respond and thus leave the dataset before the second deadline.

In Section 5, we directly investigate heterogeneity. We first demonstrate the existence of persistent types by showing that response behavior on one’s past tickets is highly predictive of response behavior on one’s current ticket (Figure 5). To explore the nature of this heterogeneity and the resulting selection, we then estimate a mixture model in which types differ in their hazard rates. Specifically, we focus on six time periods following the receipt of a ticket (i.e., we divide days 0–135 into six natural periods), and we represent each type as a vector of (period×regime)-specific hazard rates. We then estimate each type’s hazard rates jointly with the population distribution of types, allowing for up to four types.

Our estimates highlight two ways in which failing to account for heterogeneity can yield misleading conclusions. First, failing to account for selection yields biased estimates. Indeed, the estimated heterogeneity in our data implies selection effects so strong that they sometimes cause the aggregate hazard rate to decrease from one period to the next in spite of an *increase* in each of the type-specific hazard rates (Table 6). In other words, failing to account for selection sometimes reverses one’s conclusion.

Second, our estimates highlight the importance of incidence analysis: without it, one may focus on the wrong policy interventions. In particular, our estimates suggest that deadlines and reminders mostly impact types with high hazard rates, helping them to respond in a very timely way instead of in a mostly timely way, thus having no lasting impact on response rates (Figure 6). In contrast, the types with low hazard rates exhibit significant delays, react little to deadlines and reminders, and incur large costs by the time their case arrives in court (after the third deadline). Hence, our analysis suggests that if one wants to induce overall higher response rates, the focus should be less on deadlines and reminders, and more on identifying the non-responsive types and tackling the difficult question of how to get them to act. We provide some thoughts in our concluding discussion in Section 6.

Task completion was not a central topic of interest in economics until researchers started exploring sources of *suboptimal* delay.² Initial theoretical work studied procrastination on

²Our discussion here focuses on one-time tasks. Other papers focus on recurring tasks such as taking one’s daily medicine or sticking to one’s exercise regime. The latter tasks likely involve a different (though overlapping) set of mechanisms.

tasks due to present bias (Akerlof (1991); O’Donoghue and Rabin (1999); O’Donoghue and Rabin (2001)), followed by some experimental tests for the existence of procrastination (e.g., Ariely and Wertenbroch (2002)). Later work focused on the role of mental “goals” in inducing task completion (Koch and Nafziger (2011); Hsiaw (2013)), and the role of forgetting of (or inattention to) tasks (Taubinsky (2014); also see Ericson (forthcoming) who considers both procrastination and forgetting). There is experimental work documenting the existence of forgetting (Ericson (2011); Taubinsky (2014); Tasoff and Letzler (2014), who also study whether task completion can be enhanced via reminders), including field experiments that investigate the impact of reminders (e.g., Schoar and Cadena (2011)) or “planning prompts” (Nickerson and Rogers (2010); Milkman, Beshears, Choi, Laibson, and Madrian (2011)).

Relative to the latter work, our heterogeneity results suggest a word of caution, and in particular raise the question of who benefits—and, equally importantly, who does not—from interventions designed to improve task performance. Moreover, this word of caution extends to an expanding list of disciplines. Legal scholars, for example, have recently proposed using deadlines as “nudges” (Zamir, Lewinsohn-Zamir, and Ritov (2015))—using them “as a regulatory means, compared to alternative tools, such as default rules and required choices.” To the best of our knowledge, our paper is the first to focus on the *incidence* of behavioral reactions to reminders and deadlines. Our findings highlight the importance of such analysis, and the insufficiency of aggregate response evidence (typical of small-scale studies), for planning many types of real-world interventions.

Much of our contribution comes from the sheer magnitude of our dataset. Indeed, we provide graphical depictions of behavior without confidence bands (as in Figure 1) because those confidence bands are mostly indistinguishable from the depicted point estimates, and essentially any visible difference in our figures is statistically significant (see Appendix 2). Having many observations is particularly important if one wants to follow a population for a long time. For example, by day 100, roughly 85% of ticket recipients have responded and left the sample, but there are still nearly 500,000 outstanding tickets per regime that can be used to estimate hazard rates. Finally, the large dataset is not merely valuable, but is indeed crucial, when we study heterogeneity and the incidence of behavioral reactions—something that much of the previous literature simply cannot do. For example, we use our estimated

three-type model to assign each plate to one of the three types, and estimate daily hazard rates for each type (Figure 6). By the second deadline, only 1.5% of the high-type plates remain in the sample, but they still have roughly 5,000 outstanding tickets per regime that we use to estimate daily hazard rates.

2 An Organizing Model

A person receives a parking ticket on day $d = 0$ with a fine amount f . On each day $d \in \{0, 1, 2, \dots\}$ she decides whether or not to pay the ticket. If she has not paid by an initial deadline d_1 , a late penalty a_1 is imposed. Similarly, if she misses a second deadline $d_2 > d_1$, a second late penalty a_2 is imposed. Finally, if she misses a third deadline $d_3 > d_2$, she incurs an additional (exogenous) utility cost z . As we discuss in Section 3, in our field context the situation becomes complicated after the third deadline; z is meant to capture the third late penalty as well as all expected effort and other costs that might occur in the further future.

2.1 A Standard Search Model

A natural explanation for why people do not pay their tickets immediately is that they wait for a convenient (low-cost) time. To capture this intuition, we use a version of the (now standard) McCall (1970) job-search model.

We assume that paying a ticket takes a little effort and time, whose cost varies from day to day—e.g., due to daily variation in the opportunity cost of one’s time. Let c_d denote the realized effort cost on day d , drawn i.i.d. from some distribution F . A person knows her F . Each day d she learns the realization c_d and then decides whether to pay the ticket today at a known cost c_d or wait and potentially be able to pay the ticket at a lower cost in the future. The incentive to wait depends on a tradeoff between the chance to find a lower-cost day (which depends on F) and the risk of getting stuck with higher costs (which depends on the structure of deadlines and penalties).

For simplicity, we assume that the person seeks to minimize her expected total cost C , which depends on payment day \bar{d} as follows:

$$C(\bar{d}) = \begin{cases} f + c_{\bar{d}} & \text{if } \bar{d} \leq d_1 \\ f + a_1 + c_{\bar{d}} & \text{if } \bar{d} \in (d_1, d_2] \\ f + a_1 + a_2 + c_{\bar{d}} & \text{if } \bar{d} \in (d_2, d_3] \\ f + a_1 + a_2 + z & \text{if } \bar{d} > d_3. \end{cases}$$

Proposition 1 describes behavior. It essentially reiterates known results from similar optimal stopping problems (e.g., Bertsekas (2005)).

Proposition 1: (1) For each $d \leq d_3$, there exists a \underline{c}_d such that the person completes the task in period d if and only if $c_d \leq \underline{c}_d$, unless either the person does not pay the ticket in period d for any c in the support of F or the person pays the ticket in period d for every c in the support of F . (2) Moreover, for all $d \notin \{d_1, d_2\}$, if $F(\underline{c}_d) \in (0, 1)$ then $\underline{c}_{d+1} > \underline{c}_d$.

Part (1) of Proposition 1 merely states that each day the person follows a cutoff rule under which she pays the ticket today if and only if today's cost is low enough, and otherwise she waits. The more substantive result is part (2), which states that, as time passes, the person becomes increasingly likely to pay the ticket (i.e., the cutoff cost increases with time). The only exception is when the person misses a deadline, in which case she might become less likely to pay on the day immediately after the deadline. Figure 2 illustrates this pattern for a set of parameters chosen to approximate our setting (see figure notes for details). Intuitively, the person faces a trade-off: she would like to pay the ticket before the next deadline (to avoid the penalty), however, she would also like to find a convenient time. Well in advance of a deadline, it is safe to wait for a future low-cost day. As that deadline approaches, however, the incentive to pay rises because the person wants to avoid either missing the deadline or being forced to incur a high cost at the deadline. Once that deadline passes, the person is now focused on the next deadline, and thus the incentive to pay can drop immediately after the deadline, but then it rises again toward the next deadline. Finally, the spikes for the second and third deadlines are progressively larger because the penalties for missing those deadlines are progressively larger.

By assuming that the person seeks to minimize the expected total cost, we assume away discounting. Indeed, under exponential discounting, it is natural to assume a discount factor

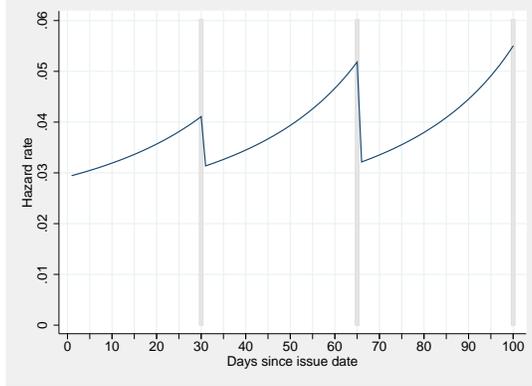


Figure 2: **Typical Pattern of Hazard Rates in Baseline Model.** Note: Hazard rates implied by baseline model under the following parameter values: $a_0 = 65$, $a_1 = 10$, $a_2 = 20$, $z = 160$, $d_1 = 30$, $d_2 = 65$, $d_3 = 100$, and $F(c) = 0.02 + c/1000$ for c between 0 and 998.

≈ 1 for daily decisions. A more relevant form of discounting is present bias (as in Laibson (1997), O’Donoghue and Rabin (1999)), as it seems plausible that people are prone to procrastinate paying their tickets. However, the direct impact of present bias in this model is merely to reduce the likelihood of acting in any particular period, without changing the qualitative pattern of behavior illustrated in Figure 2. Moreover, it is difficult to separately identify discounting parameters and the distribution of effort costs (F) without variation in the magnitude of penalties or the timing of deadlines—variation that we do not have in our data. Hence, while we believe present bias is likely playing a role in this environment, we exclude it from our organizing model.

A comparison between Figures 1 and 2 suggests that, except for the spike at the first deadline, the qualitative pattern predicted by the model does not capture the data well. The increasingly large spikes at the second and third deadlines predicted by the model are all but absent in the data, and moreover the model does not explain the dramatic decline in hazard rates between the first and second deadline. Incorporating forgetting and heterogeneity into the model can bring it closer to the data, as we discuss next.

2.2 Heterogeneity and Selection

Prior analyses of search—especially of job search—attribute declines in hazard rates over time to the existence of heterogeneity and selection. The intuition is simple. Suppose there

are two types, call them High and Low. Both behave according to the model above, but High are more prone to pay on any given day than Low—i.e., they face a different distribution of costs (F). As time since day 0 passes, a growing share of the population of people with unpaid tickets is Low. Hence, the observed hazard rate, which is a weighted average of the two type-specific hazard rates, will gradually drop.

To test for the existence of persistent types in the population, one can test whether an individual’s response behavior on some tickets helps predict response behavior on other tickets.³ We pursue this strategy to investigate heterogeneity in Section 5.

2.3 Forgetting

A second reason why hazard rates might decline over time is forgetting. Simple “i.i.d. forgetting,” wherein on each day there is some probability p that a person remembers her outstanding ticket, cannot generate declining hazard rates.⁴ We consider instead forgetting with persistence. Suppose that individuals are all of the same type (same F , etc.), but each can be in one of two internal states, call them On-the-mind and Off-the-mind (see Taubinsky (2014) for a model of this kind). If a ticket is on a person’s mind today, then it is likely to remain on her mind tomorrow. However, once the ticket slips off the mind, it is likely to remain off her mind tomorrow. In such a model, the longer a ticket is outstanding, the more likely it is to be off the mind, and thus the lower should be the observed hazard rate.

To test for off-the-mind forgetting, one can study the impact of interventions designed to put the ticket back on the mind (at least with some probability). Such interventions should lead to a jump in hazard rates and a restart of the forgetting process. In Section 4, we pursue such a test by using exogenous changes in the timing of notification letters.

³Notice that for a known type, hazard rates are the same for all tickets. Hence, in the one-type model in Section 2.1 (and Figure 2), where a person’s type is known, there is no information in past behavior to help predict current behavior. In contrast, in the two-type model described here, the later a person paid a past ticket, the more likely she is to be Low, and thus she is predicted to pay the current ticket later.

⁴If we modified the model to include an underlying cost distribution $\hat{F}(c)$ and i.i.d. probability of remembering p , it would be equivalent to the model in Section 2.1 with cost distribution $F(c) = p\hat{F}(c)$ —and thus Proposition 1 still holds.

3 New York City Parking Tickets

3.1 Data Description

Appendix 1 contains a detailed description of our data. Here we summarize the most important details. The data come from the New York City Department of Finance (DOF), which handles most incoming revenue to the city. The full dataset contains information on 20,874,688 tickets, covering (virtually) all tickets issued between June 1, 2011 and August 31, 2013.⁵ The data include ticket issue date, violation type, fine amount, issuing agency, location, and many other details. In addition, the data allow us to construct each ticket’s history of “events” through late January 2014—i.e., actions taken either by the ticket-receiving individual (e.g., making a payment or contesting the ticket) or by DOF (e.g., imposing a late penalty or sending a notification letter).

We restrict the full dataset in several ways. Most notably, we exclude: tickets issued for commercial vehicles and vehicles that are rentals or part of a fleet program (32.4% of tickets in the full dataset); two moving violations which fall under the purview of DOF (red-light and bus-lane violations caught with stationary cameras, which are subject to different rules); and several non-moving, non-parking violations (e.g., expired registration). These restrictions leave us with a *passenger dataset* that contains 11,139,375 tickets.

Next, to ensure that notification letters were received, we limit attention to vehicles with DOF’s highest address verification level, leaving us with 6,801,115 tickets. We then drop the 173 remaining tickets with addresses outside of New York, and restrict attention to the 23 most common parking violations, leaving us with 6,730,378 tickets. Finally, while tickets can be issued by many agencies, the vast majority are issued by either New York City parking ticket agents or by the New York City Police Department, and we restrict attention to those. The remaining 6,646,540 tickets comprise the *core dataset* that we analyze in this paper.

The first column of Table 1 presents descriptive statistics for the core dataset (see Appendix Table A1 for a comparison with the passenger dataset). The most common violations are for expired parking meter (36.2%), no parking due to street cleaning (26.2%), parking

⁵Due to what appears to be a data-transfer issue, we are missing a large number of tickets issued on two specific days (February 27–28, 2013).

in a general no parking zone (9.2%), and parking or standing in a general no standing zone (6.7%).⁶ The most common fine amounts are \$35 (30.1%), \$45 (23.9%), and \$115 (23.0%). The vast majority (97.2%) of tickets are issued by parking-ticket agents; indeed, excluding those issued by police officers does not affect our results.

Four payment methods are available: by mail, online, by phone, or in person (at one of five DOF Business Centers). The bottom panel of Table 1 presents the distribution of payment types for the 80.2% of tickets in the core dataset that have payments made by day 135. The most common payment methods are online (53.8%) and by mail (32.3%).

3.2 The OLD and NEW Regimes

Table 2 summarizes the timeline of key events under the OLD and NEW regimes (and also under the EXP regime, which we describe in Section 4.3). The timeline is identical for the OLD and NEW regimes except for one thing: DOF changed the timing and the content of the first notification letter. The rest of this section provides detail.

3.2.1 Timeline in the OLD Regime

Tickets are issued (on what we define as day 0) as a paper ticket that is placed, along with an orange envelope, on the windshield of the offending car. The ticket and envelope together provide information on the violation type, fine amount, the four possible ways to pay, information on how to contest the violation, the (first) due date of day 30, and the (first) late-payment penalty amount of \$10. They also mention that failure to respond may result in additional penalties and a default judgment being entered, which in turn could mean the vehicle might be towed. Appendix 13 contains sample tickets and the relevant part of the envelope, as well as samples for all notification letters described below.

While the stated due date is day 30, in practice DOF’s computer system runs a batch job every Saturday, and assigns a \$10 late penalty to any outstanding ticket for which the due date passed during the preceding week ending on Thursday. In addition, DOF also mails a

⁶In Table 1, the expired-meter category aggregates three separate violations for different types of meters, and the authorized-vehicles-only category aggregates two separate violations. Hence, there are 20 rather than 23 violation types in Table 1.

notification letter to the plate owner (OLD letter 1) on the subsequent Tuesday (which would be day 35–41). This letter, labelled “NOTICE OF OUTSTANDING VIOLATION,” shows an updated balance due that includes the \$10 late penalty. It also provides a new due date, the Monday that is exactly 27 days after that Tuesday (day 62–68), and states that failure to respond in time will result in an additional late penalty of \$20 and can lead to a default judgment entry, which would then permit actions such as garnisheeing the owner’s wages, towing the owner’s vehicles, and preventing renewal of motor vehicle registrations. Note that the list of possible actions following a default judgment entry has expanded relative to the language on the original ticket/envelope—and will keep expanding in later letters, including increasingly severe possible actions—a point we return to below. Finally, like the original ticket/envelope, the letter also describes the payment and contest methods.⁷

If no response is received by the first Saturday after the second deadline, the second late penalty is applied and in addition DOF mails a second notification letter (letter 2) on the subsequent Tuesday (day 70–76). Letter 2, labelled “NOTICE OF IMPENDING DEFAULT JUDGMENT,” shows an updated balance that includes the second late penalty, and provides yet a new due date, which is the Friday that is exactly 31 days after that Tuesday (day 101–107). The letter warns that failure to respond will lead to a default judgment entry against the owner, which would then permit actions such as reporting the debt to credit agencies, towing the owner’s vehicles, garnisheeing the owner’s wages and/or seizing assets including real estate and bank accounts, and preventing renewal of motor vehicle registrations. However, the letter does not explicitly mention the amount—\$30—of the impending third late penalty.

If no response is received by the first Saturday after the third deadline, a third late penalty of \$30 is applied. In addition, DOF sends a third notification letter (letter 3) on the subsequent Tuesday (day 105–111), stating that it has entered a default judgment, and that the owner is now subject to immediate judgment enforcement procedures. The due date on letter 3 is “IMMEDIATELY” and the letter indicates a range of severe punishments

⁷Each notification letter included four pages, the last three of which describe the details of possible ways to pay and to contest (DOF is required by state law to have certain language in the notices). Those three pages are the same for all letters, except for one place on page 4 that lists the details of the specific violation (violation type, fine amount, etc.). In Appendix 13, we provide those three pages only for OLD letter 1.

such as “Assigning your judgment debt to a Collection Agency,” “Seizing any motor vehicle registered to you and selling it at auction,” and “Seizing your non-exempt personal property.”

3.2.2 Change to the NEW Regime

Under the OLD regime, OLD letter 1 was sent shortly after the first deadline (on day 35–41). Beginning with tickets issued on June 18, 2012, DOF moved this letter to before the first deadline. Specifically, if there is no response by day 17, then a letter (NEW letter 1) is generated on day 18, mailed on day 19, and (most likely) received on day 20. The exceptions are tickets issued on Tuesday or Wednesday, for which day 18 occurs on a weekend. For these tickets, the letter is generated on the subsequent Monday (day 19 or 20), mailed on Tuesday (day 20 or 21), and (most likely) received on Wednesday (day 21 or 22).

In addition, while most of the content of NEW letter 1 is identical to that of OLD letter 1, DOF made three substantive changes. First, the label is changed from “NOTICE OF OUTSTANDING VIOLATION” to “PRE-PENALTY NOTICE OF UNPAID VIOLATION.” Second, instead of stating the second deadline (day 62–68) and the second penalty (\$20), NEW letter 1 states the first deadline (day 30) and the first penalty (\$10). Finally, unlike OLD letter 1, which mentions the possibility of a default judgment entry and uses a bold font to highlight various future penalties, NEW letter 1 does not mention the possibility of a default judgment entry and does not contain any bold font.

We highlight two further differences between the two regimes. First, prior to February 1, 2012, DOF had a settlement program in which, if an owner initiated a contest, the owner was automatically offered a fine reduction if they accepted a settlement instead of continuing with the contest. The offered fine reduction varied by ticket type, ranging from \$10 to \$25. This program was abolished because DOF had come to believe that drivers were gaming the program. While this program was available during much of the OLD regime but never during the NEW regime, its elimination had little effect on our results (see Section 4.1).

Second, as described above, communications from DOF contain incomplete information about the deadlines and the possible consequences of missing them. Consider a driver without prior knowledge of the timeline of events. Under the OLD regime, she is initially (in the ticket/envelope) told only about the first deadline and its associated \$10 penalty. Only

after missing that deadline is she told (in OLD letter 1) about the second deadline and its \$20 penalty. Similarly, only after missing the second deadline is she told (in letter 2) about the third deadline, but only after missing it is she told (in letter 3) the amount (\$30) of the third penalty—after it has already been imposed. In addition, from the start, the driver repeatedly receives hints at the possibility of a default judgment entry and ensuing enforcement actions, but is not told when that will occur until letter 3—when a default judgment has now been entered. Under the NEW regime, things are much the same, except that now no notification letter is sent between the first and second deadline, and thus the driver is not told about the existence of the second deadline and its \$20 penalty until after that deadline has passed and the penalty has been imposed. We return to this issue when we motivate our field experiment, which was designed in part to test whether differences in information might drive some of the observed behavioral differences across the two regimes. As we discuss there, this issue seems not to matter.

3.2.3 Descriptive Statistics

The OLD regime applies to tickets issued between June 1, 2011 and June 17, 2012. The NEW regime applies to tickets issued between June 18, 2012 and July 12, 2013 and between August 17, 2013 and August 31, 2013. (Tickets issued between July 13, 2013 and August 16, 2013 were part of the EXP regime, which we discuss in Section 4.3.) As the second and third columns of Table 1 show, the distributions of violation type, ticket amount, ticket issuer, and payment type are all similar across the regimes. The main difference is that as time passes and we move from OLD to NEW regime, there is a modest shift from in-person and mail payments to online and phone payments (a shift that continues into the EXP regime).

4 Evidence of Forgetting

4.1 First Responses

Drivers in our setup face a two-dimensional decision: they choose not only *when* to respond to a ticket, but also *how* to respond—they can pay or contest. We focus on the *when* question—

that is, we study the timing of first responses of any type, pooling together payments and contests (as in Figure 1).

Table 3 describes the when and how of first responses. Timing is measured in days since issue date and can take values from 0 to 135, or a no-response indicator if no response is observed by day 135. First response can be a payment (including partial) or a contest, whichever happens sooner.⁸ We count settlements through the settlement program as a contest. The top section of the table shows the distribution of first responses for all tickets.

The cumulative response rate through day 135 is roughly the same across the OLD, NEW, and EXP regimes—90.1%, 89.3%, and 89.5%. However, there are differences in the type of first response, with more contests and fewer payments in the OLD regime than in the NEW and EXP regimes. This difference is primarily due to the elimination of the settlement program on February 1, 2012. To show this, Table 3 also presents first responses for the subset of tickets under the OLD regime that were issued on February 1, 2012 or later, which we label the *OLD-post* regime. Cumulative behavior through day 135 is much the same (in fact, slightly closer to NEW and EXP), and moreover the composition of first responses is similar under the OLD-post, NEW, and EXP regimes.

Our analysis of first responses pooled across response types creates two potential worries that Table 3 helps to alleviate. First, within a regime, we may be missing interesting patterns in the type of response over time. Second, and more important, when comparing regimes, we may be focusing on the wrong question—the timing question—if the primary impact of the regime change is on the type of first response. Our sense is that the regime changes are unlikely to have much impact on the type of first response, and indeed we see no evidence of this in the data. Specifically, the bottom panel of Table 3 reports first-response behavior across different time intervals as well as the composition of first responses within an interval. The regimes clearly differ in the timing of pooled first responses—i.e., the NEW and EXP regimes have more first responses prior to deadline 1 than the OLD regime, and fewer first responses between deadlines 1 and 2. This is the pattern highlighted in Figure 1 that we

⁸The number of payments by day 135 differs in Tables 1 and 3 because the former reflects all payments by day 135 while the latter only reflects payments that are also a first response. For a small fraction of contests, we observe only the resolution date, which might come zero or more days after the contest (see Appendix 1 and Appendix 10 for details).

explore in detail in Section 4.2. However, the data exhibit two additional features that together alleviate our concerns above. First, the timing of *pooled* first responses in the OLD-post regime is roughly the same as that in the OLD regime, and thus the elimination of the settlement program appears to have altered the type of first response without altering the timing of pooled first responses. Second, the within-interval *composition* of first responses in the OLD-post regime is roughly the same as that in the NEW and EXP regimes, especially prior to deadline 2, and thus the regime shifts from OLD-post to NEW to EXP appear not to have altered much the type of first response.

4.2 Behavior under the OLD vs. NEW Regimes

We analyze first-response hazard rates using survival analysis. Each ticket is a single observation, and its first response is viewed as a failure event. We estimate daily hazard rates by dividing, for each of days 0–135, the number of first responses on that day by the number of outstanding tickets, i.e., tickets with no first response before that day.

As described in Section 1, Figure 1 depicts estimated hazard rates in the OLD and NEW regimes.⁹ Prior to day 20, behavior is roughly the same under the two regimes, as expected given there is not yet any differential treatment between regimes. Then, when NEW letter 1 hits in the NEW regime, hazard rates are increasingly larger and more obviously spike at day 30 (deadline 1) relative to hazard rates in the OLD regime. Analogously, from roughly day 40, when OLD letter 1 hits in the OLD regime, hazard rates increase relative to hazard rates in the NEW regime. After the second deadline hazard rates quickly converge, and after the arrival of the second letter (sent on days 70–76), behavior is as similar as it was prior to day 20. (Appendix 3 reproduces Figure 1 using the OLD-post regime in place of the OLD regime, and shows that the small OLD-NEW differences before day 20 and after day 77 essentially disappear, suggesting they reflect the discontinued settlement program.)

In Figure 1, different tickets involve different timing of deadlines and letters depending on the day of the week a ticket is issued. To address this issue, and as a means to quantify some of the differences seen in Figure 1, we analyze behavior across six natural “periods” (we also use these periods when we estimate a mixture model in Section 5.2).

⁹Appendix 2 reproduces the major figures with 95% confidence bands.

Period 1: from day 0 to the day NEW letter 1 is sent (day 20 and 21 for tickets issued on Wednesday and Tuesday, respectively; day 19 otherwise)

Period 2: from the day after NEW letter 1 is sent to deadline 1 (day 30 for everyone)

Period 3: from day 31 to the day OLD letter 1 is sent (day 35–41)

Period 4: from the day after OLD letter 1 is sent to deadline 2 (day 62–68)

Period 5: from the day after deadline 2 to the day letter 2 is sent (day 70–76)

Period 6: from the day after letter 2 is sent onward

Because these periods have different lengths, we focus on average daily hazard rates within each period. Table 4 presents estimated average daily hazard rates and cumulative response rates (see Appendix 4 for details of how Table 4 is created). As in Figure 1, response rates prior to day 20 (in period 1) are roughly the same across the two regimes, but then NEW letter 1 leads to a dramatic increase in response rates relative to the OLD regime, both before and for a while after the first deadline—average daily hazard rates are 30% higher in period 2, and 45% higher in period 3. However, once the first letter is sent in the OLD regime, response rates become higher than those in the NEW regime, both before and for a while after the second deadline—average daily hazard rates are 40% higher in period 4, and 45% higher in period 5. Like Figure 1, Table 4 shows that the cumulative response rate by the time letter 2 is sent is roughly the same under the NEW vs. OLD regimes.

An alternative way to address the issue that tickets differ in the timing of deadlines and letters is to redefine day 0 to be the date of deadline 2. Under this approach, OLD letter 1 is sent on day -27 for all tickets, deadline 2 is day 0 for all tickets, letter 2 is sent on day 8 for all tickets, deadline 3 is day 39 for all tickets, and letter 3 is sent on day 43 for all tickets. Figure 3 depicts hazard rates estimated with this alternative definition of days. We see the same general pattern, and moreover we now see a noticeable reaction immediately after each letter is sent, providing further evidence of forgetting.¹⁰

¹⁰Figure 3 reveals more noticeable spikes in hazard rates at deadlines 2 and 3, although these spikes are much smaller than the spike at deadline 1. It also exhibits a clear weekly cycle in hazard rates, with lower response rates on weekends. See Appendix 5 for hazard rates estimated by issue day of week.

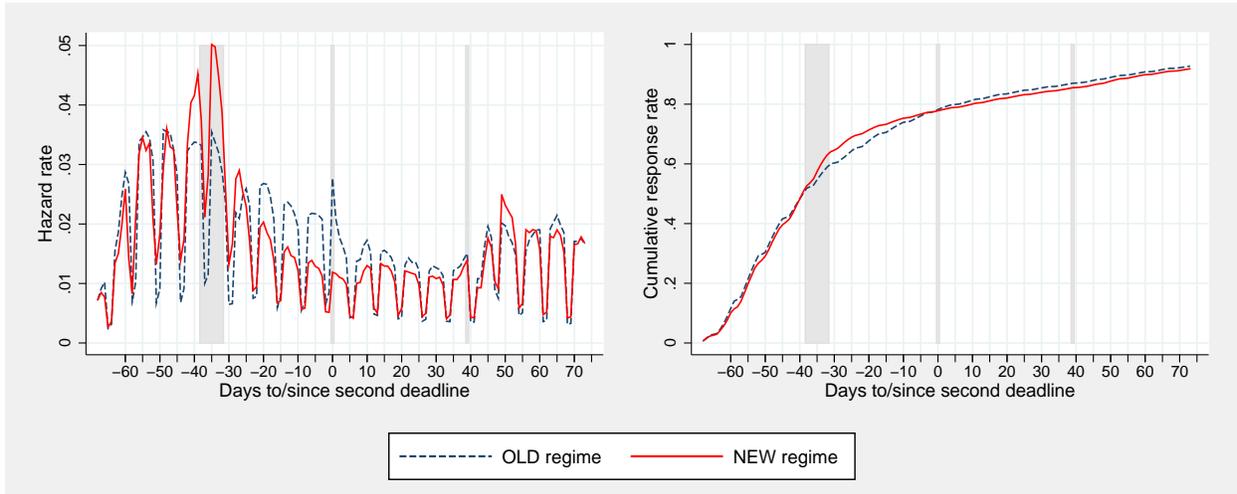


Figure 3: **Hazard Rates and Cumulative Response Rates in OLD vs. NEW Regimes.** Note: Day 0 is the day of the second deadline, which is a Monday for all tickets. All tickets have the third deadline at day 39 (a Friday), and the first deadline at days -38 to -32 , indicated by shaded areas. First notification letter is received on day -26 (OLD) vs. about day -46 (NEW).

Finally, while Section 5 focuses on unobserved heterogeneity, here we assess the robustness of our conclusions about forgetting to *observed* heterogeneity. Specifically, within each regime, we estimate daily hazard rates separately for each of (i) the six most-common violation types, (ii) the six most-common fine amounts, and (iii) the two issuing agencies (Appendix 6 contains the figures). From this analysis, we draw two conclusions. First, while there are noticeable differences across sub-groups, there is nothing systematic that relates naturally to some underlying mechanism (e.g., while hazard rates are different for different fine amounts, higher fines are associated with neither higher nor lower hazard rates). Second, within each sub-group, the qualitative comparison between the OLD vs. NEW regimes is essentially the same—that is, our results appear robust to observed heterogeneity in violation type, fine amount, and ticketing agency.

4.3 A Field Experiment (the EXP Regime)

After some initial comparisons of behavior under the OLD vs. NEW regimes, DOF agreed to conduct a field experiment. The experiment included random variation along three dimensions: (i) NEW letter 1 might or might not include additional information (described

below), (ii) NEW letter 1 might or might not include some “scary” language (also described below), and (iii) there might or might not be a new reminder letter between the first and second deadlines (which we label EXP letter 1.5).¹¹ Hence, there are eight experimental cells as described in Table 5.

This design addresses three issues of interpretation in the comparison of the OLD vs. NEW regimes. First, as discussed in Section 3.2.2, drivers learn the schedule of deadlines and penalties in a piecewise fashion, and there are differences in this information across regimes. To explore whether these differences in information drive some of the differences in behavior between the OLD vs. NEW regimes, some drivers—those in the *info* and *info scary* treatments—received a modified version of NEW letter 1 that includes a box that lists the full set of (individualized) deadlines and penalties. For instance, for a ticket issued on Monday, July 15, 2013, this box would read as follows:

AMOUNT DUE IF PAID BY 8/14/13:	\$65
AMOUNT DUE IF PAID BY 9/16/13:	\$75 (INCLUDES \$10 PENALTY FOR LATE PAYMENT)
AMOUNT DUE IF PAID BY 10/25/13:	\$95 (INCLUDES \$30 PENALTY FOR LATE PAYMENT)
AMOUNT DUE IF PAID AFTER 10/25/13:	\$125 (INCLUDES \$60 PENALTY FOR LATE PAYMENT)

If no payment is received by 11/1/13, Finance may boot or tow your vehicle.

Second, the language used in NEW letter 1 is different from that used in OLD letter 1. In particular, OLD letter 1 mentions the possibility of a default judgment entry and the associated actions, and moreover uses a bold font to highlight the various future penalties. In contrast, NEW letter 1 does not mention the possibility of a default judgment entry, and does not contain any bold font. To investigate the impact of such language differences, some drivers—those in the *scary* and *info scary* treatments—received a modified version of NEW letter 1 that contains arguably scarier (or more forceful) language. Specifically, the letter had the following header in large bold-faced letters:

¹¹Varying penalty amounts and deadlines might make it possible to identify discounting and possibly present bias, but such interventions were not feasible.

WARNING: PENALTY APPROACHING
DON'T MISS THE DEADLINE

In addition, NEW letters *1i*, *1s*, and *1is* all mention that failure to respond might result in one's vehicle being booted or towed, and in NEW letters *1s* and *1is* this is mentioned in a larger font size.

Third, in our comparison of the OLD vs. NEW regimes, we interpreted the higher response rates in the OLD regime between OLD letter 1 and the second deadline as further evidence of forgetting, but this comparison is not clean due to different treatment prior to OLD letter 1. To get a cleaner test of the impact of a letter between the first and second deadlines, some drivers, chosen independently of the NEW letter 1 groups, received a new reminder letter (EXP letter 1.5) between the two deadlines. Specifically, if there is no response by day 45, then a letter is generated on day 46, mailed on day 47, and (most likely) received on day 48, except for tickets issued on Tuesday or Wednesday, for which day 46 occurs on a weekend and the letters are generated on the subsequent Monday. The content of this letter is identical to that in NEW letter *1i*, except that (i) the first amount due in the information box is deleted (since it is no longer relevant) and (ii) the letter is labeled "NOTICE OF OUTSTANDING VIOLATION."

The experimental (EXP) regime applied to all tickets issued from July 13, 2013 through August 16, 2013. For tickets issued during these five weeks, if a NEW letter 1 was triggered, it was randomly assigned to one of the four NEW letter 1's according to the probabilities in Table 5. The EXP letter 1.5 treatment applied only for tickets issued July 22, 2013 through August 10, 2013 (and implemented by generating EXP letter 1.5 over three business weeks in September 2013). For tickets issued during this period, if an EXP letter 1.5 was triggered, with 50% chance an EXP letter 1.5 was sent (regardless of which NEW letter 1 was sent).¹²

Table 1 provides descriptive statistics for the EXP regime. In general, these statistics look roughly the same as both the OLD and NEW regimes. Table 3 presents an overview of first responses in the EXP regime (see the discussion in Section 4.1).

¹²Each randomization was done by ordering plates alphanumerically and then assigning plates to treatments via a pre-set pattern. For plates that received multiple tickets in the EXP regime, it was possible to receive different treatments for the different tickets. Our results in Section 4.4 are unchanged if we consider only plates that received exactly one ticket in the EXP regime (see Appendix 7.1).

4.4 Behavior under the EXP Regime

We analyze daily hazard rates in each of the eight experimental cells. However, because randomization occurred only when letters were generated, and not at the time tickets were issued, we create the eight experimental cells by performing an ex post random assignment for all tickets with a response prior to the day-18 or day-46 letter generation.¹³ Specifically, for each ticket with a response prior to day 18 (that did not receive any NEW letter 1) we ex post randomly assign it into one of the four NEW letter 1 treatments, using the same probabilities as in Table 5. Similarly, for each ticket issued between July 22, 2013 through August 10, 2013 for which there is a response prior to day 46, we ex post randomly assign it into one of the two EXP letter 1.5 letter treatments. Finally, we assign all tickets issued outside of July 22, 2013 through August 10, 2013 to the no-EXP-letter-1.5 treatment, since none of them could have received an EXP letter 1.5.¹⁴ After applying this ex post randomization, the number of observations in the four cells without EXP letter 1.5 are 38,016 (1), 76,510 (1*i*), 38,321 (1*s*), and 38,130 (1*is*), and the number of observations in the four cells with EXP letter 1.5 are 16,053 (1), 32,133 (1*i*), 15,854 (1*s*), and 16,072 (1*is*).¹⁵

Figure 4a depicts hazard rates for the four experimental cells assigned not to receive an EXP letter 1.5. It reveals that the four versions of NEW letter 1 lead to almost identical hazard rates. This result suggests that neither the large differences in behavior between the OLD vs. NEW regimes nor the weak response to the third deadline in both regimes are driven by differences in information or language (recall that the weak response to the third deadline in both regimes could have resulted from individuals not having been informed about the \$30 penalty; however, informing them about all penalties and deadlines in the *info* and *info scary* treatments had little impact).¹⁶

¹³As shorthand, we refer to the day on which NEW letter 1 is triggered as “day 18,” even though for some tickets it is day 19 or day 20 (see Section 3.2.2). For EXP letter 1.5, we analogously use “day 46.”

¹⁴Our results are robust to dropping tickets outside of July 22 to August 10, 2013 (see Appendix 7.2).

¹⁵An alternative approach would be to conduct a single hazard rate estimation prior to receipt of NEW letter 1, then four hazard rate estimations between NEW letter 1 and EXP letter 1.5, and finally eight hazard rate estimations after EXP letter 1.5. We chose not to pursue this approach because of difficulties associated with the split dates varying depending on the day of the week on which a ticket is issued.

¹⁶A close inspection of the confidence intervals in Appendix 2 reveals that the *scary* and *info scary* treatments led to a statistically significant but very small increase in hazard rates relative to the *baseline* and *info* treatments. Because NEW letter 1 contains less scary language than OLD letter 1, if anything our comparison of the OLD vs. NEW regimes slightly *understates* the impact of NEW letter 1.

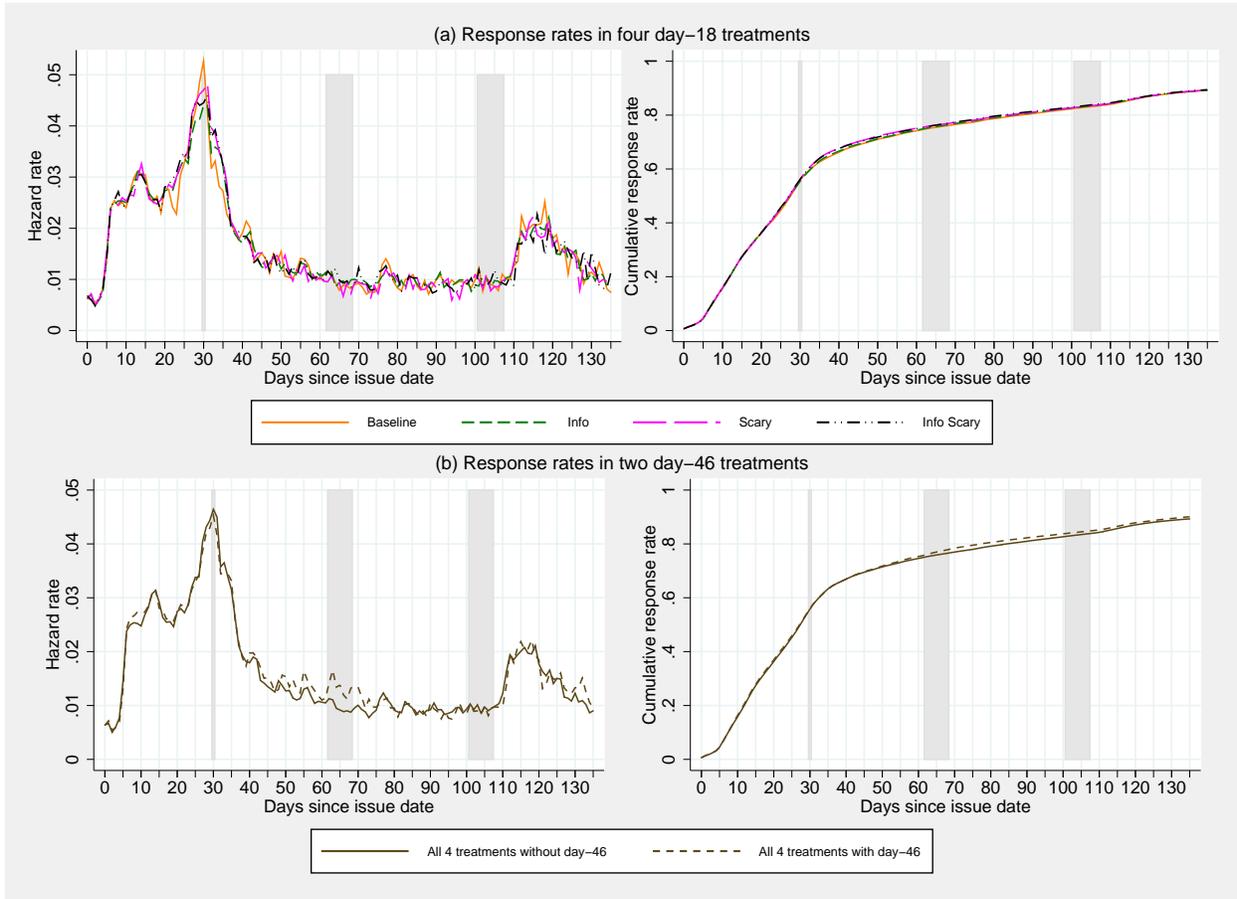


Figure 4: **Hazard Rates and Cumulative Response Rates in EXP Regime.** Note: Figure 4a includes only the four experimental cells in which EXP letter 1.5 is not sent. Figure 4b pools the four experimental cells assigned not to receive, and to receive, an EXP letter 1.5.

Figure 4b, in which the four NEW letter 1 treatments are pooled together, reveals that EXP letter 1.5 has a noticeable impact, further supporting the role of off-the-mind forgetting. Even after getting a letter shortly after day 18, getting a second letter shortly after day 46 increases response rates. However, this effect is much smaller than the impact of NEW letter 1 under the NEW and EXP regimes, and of OLD letter 1 under the OLD regime. These differences in magnitude suggest the existence of heterogeneous types and selection: with heterogeneity, under the EXP regime the distribution of types that receives EXP letter 1.5 is different from the distribution that receives NEW letter 1 (and it is also different from the distribution receiving OLD letter 1 in the OLD regime) and thus the significant differences in responses. We turn our attention to such selection shortly.

The differential impact of the first and second letters raises two further issues. First,

while we interpret the impact of the first letter as a reminder, it is possible that, in some instances, the first letter is actually the first time the owner learns of a ticket—perhaps the ticket was received by someone who borrowed the car, or perhaps the ticket somehow fell off the car before the owner noticed it. However, our type results in Section 5 would seem to contradict this story, because it would require that the good types are the types prone not to know about the ticket, and this seems implausible. Second, there might be something special about the first letter that one receives—perhaps it reveals to owners that DOF knows where they live, or that the ticket really does need to be paid. However, such effects would primarily apply for first offenders, and our results look much the same for repeat offenders—as seen in Figure 5 in Section 5, or by comparing Tables A4 and A5 in Appendix 8.1.

5 Unobserved Heterogeneity and Selection

5.1 Evidence of Persistent Types

To demonstrate the existence of persistent types, we show that response behavior on one’s past tickets is highly predictive of response behavior on one’s current ticket. We use a relatively simple approach. First, we identify all plates that received exactly three tickets under the OLD regime, and divide them into four groups: (i) plates whose first two tickets are both responded to by day 30, (ii) plates whose first ticket is, but second ticket is not, responded to by day 30, (iii) plates whose second ticket is, but first ticket is not, responded to by day 30, and (iv) plates whose first two tickets are neither responded to by day 30. Then, for each of these four groups, we estimate daily hazard rates for each plate’s third ticket. We carry out an identical exercise for all plates that received exactly three tickets under the NEW regime.¹⁷

Figure 5 depicts these hazard rates. We see clear evidence of persistent types. Third-ticket hazard rates for plates in group (i) are roughly twice the average daily hazard rates in Figure 1, while those in group (iv) are less than half the rates in Figure 1, and those in groups (ii) and (iii) are in between. Moreover, the allocation of plates into the four groups

¹⁷In the OLD regime, the number of observations in each group are 56,035, 19,872, 20,429, and 41,559. In the NEW regime, the number of observations in each group are 55,783, 17,510, 17,166, and 35,111.

itself provides further evidence of persistent types. Specifically, the vast majority of plates fall into either group (i) (40.6% in the OLD regime and 44.4% in the NEW regime) or group (iv) (30.1% and 28.0%). These allocations are inconsistent with every ticket having the same independent probability to respond by day 30. For instance, seeing 40.6% in group (i) implies that the probability of responding by day 30 is 63.7% (because $0.637^2 = 0.406$), but that in turn would mean that we should see only 13.2% in group (iv) (because $(1 - 0.637)^2 = 0.132$).

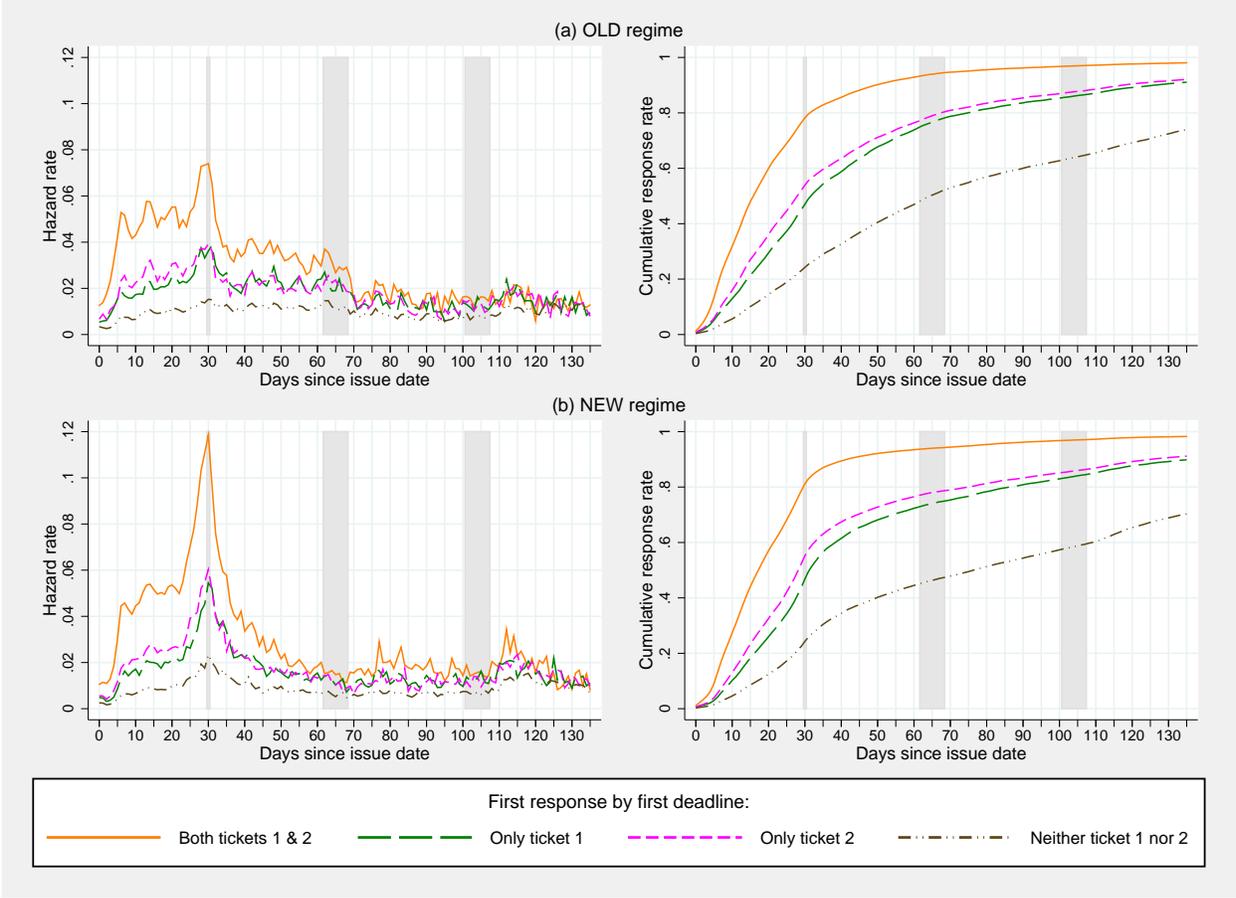


Figure 5: **Past Response Behavior Predicts Current Response Behavior.** Note: Daily hazard rates and cumulative response rates split by the ticket recipient’s response behavior on two prior tickets. Based on recipients of exactly three tickets within a regime.

5.2 A Mixture Model of Unobserved Types

Having demonstrated the existence of persistent types, we next estimate a mixture model of unobserved types. We impose no structure on hazard rates over time. However, for

tractability, instead of considering a model at the daily frequency, we conduct this analysis in terms of the six periods introduced in Section 4.2.

Suppose there is a discrete set K of types in the population, and let π_k denote the proportion of the population that is type $k \in K$, where $\sum_{k \in K} \pi_k = 1$. Each type k is characterized by hazard rates $(p_1^k, p_2^k, p_3^k, p_4^k, p_5^k)$. The hazard rate p_t^k is the probability that the person responds to a ticket in period t conditional on not having responded prior to period t . The hazard rate in the last (open-ended) period 6 is by definition equal to 1.

We permit hazard rates to depend on the regime γ , and thus write $p_t^k(\gamma)$ for each $t \in \{1, 2, 3, 4, 5\}$ and $k \in K$. The impact of regime on the p_t^k 's will be a measure of the impact of the reminder letters—e.g., the extent to which $p_2^k(\text{NEW}) > p_2^k(\text{OLD})$ is a measure of the immediate impact of NEW letter 1 in the NEW regime (relative to no letter in the end of period 1 in the OLD regime), and the extent to which $p_4^k(\text{OLD}) > p_4^k(\text{NEW})$ is a measure of the immediate impact of OLD letter 1 in the OLD regime (relative to no letter in the end of period 3 in the NEW regime).¹⁸

For plate i , we can write the observed behavior as a vector

$$d^i \equiv (J_i, m_1^i, \gamma_1^i, m_2^i, \gamma_2^i, \dots, m_{J_i}^i, \gamma_{J_i}^i).$$

where J_i is the total number of tickets received by plate i , $m_j^i \in \{1, 2, 3, 4, 5, 6\}$ is the period in which plate i 's owner responded to ticket j , and γ_j^i is the regime that applies to ticket j for plate i . Then, conditional on receiving J_i tickets, the likelihood that type k would generate observed behavior d^i is

$$\ell_k(d^i) = \prod_{j=1}^{J_i} \left([p_1^k(\gamma_j^i)]^{I\{m_j^i=1\}} \prod_{t=2}^6 \left(\left[\prod_{t'=1}^{t-1} (1 - p_{t'}^k(\gamma_j^i)) \right] p_t^k(\gamma_j^i) \right)^{I\{m_j^i=t\}} \right)$$

where I is the identity function, and thus the likelihood that plate i generates observation

¹⁸The latter statement is true for a known type. In practice, however, given that our estimation will permit only a small number of types, there is likely some residual heterogeneity within an estimated type, and thus part of the difference between $p_4^k(\text{OLD})$ and $p_4^k(\text{NEW})$ might be due to some differential selection in periods 2 and 3.

d^i is

$$\ell_i(d^i) = \sum_{k=1}^K \pi_k \ell_k(d^i).$$

Finally, assuming that the number of tickets received J_i is independent of one’s type k , the sample log-likelihood can be written as¹⁹

$$\log \mathcal{L} = \sum_i \log \ell_i(d^i).$$

This model makes several simplifying assumptions. It assumes that (i) the population distribution of types π_k is the same for each regime γ , (ii) the number of tickets received J_i is independent of one’s type k , and (iii) within a type, the $p_t^k(\gamma)$ ’s are the same for all tickets received under regime γ (this assumption rules out “learning” in the sense that one’s response experience on prior tickets does not change one’s response behavior on the current ticket, as well as any other form of interaction across tickets). These assumptions are primarily made for practical reasons, both to reduce computational burden and because there is a limit to how much heterogeneity we can identify. In Appendix 8.1, we assess these assumptions, and report evidence that, while not fully accurate, they seem reasonable for our purposes.

In our main estimation, we use all plates that received $J \in \{3, 4, \dots, 12\}$ tickets across the OLD and NEW regimes combined—657,890 plates that received 3,366,145 tickets.²⁰ Before estimating the model, for each plate we take one randomly chosen ticket and put it in a holdout sample, which we will use in Section 5.3 below. Using the remaining 2,708,255 tickets for the 657,890 plates, we estimate the mixture model above for $K \in \{1, 2, 3, 4\}$.

Table 6 reports, for each $K \in \{1, 2, 3, 4\}$, the estimated average daily hazard rates for each type in each period, along with the estimated proportion of each type.²¹ The estimated model yields several interesting conclusions. The discussion below focuses on the $K = 3$

¹⁹The actual sample log-likelihood is $\sum_i \log (\Pr(J_i) \ell_i(d^i))$, but as long as $\Pr(J_i)$ is independent of type k , it does not impact the estimation.

²⁰We do not use data from the EXP regime in estimating this model because regime-specific hazard rates are identified from plates that have multiple tickets within a regime, and there are very few plates that receive multiple tickets within any one cell in the EXP regime. We do study response behavior for the typed plates in the EXP regime in Section 5.3 below.

²¹As described above, our estimation technique yields per-period hazard rates (these are reported in Appendix Table A7). For interpretation, we convert each per-period hazard rate into an average daily hazard rate using the average number of days in each period, and use the delta method to convert the standard errors. Details of this transformation are available in Appendix 8.2.

model, but similar messages emerge from the $K = 2$ or $K = 4$ models. In the $K = 3$ model, we refer to the first type with the highest hazard rate in all periods as the *high type*, the third type with the lowest hazard rate in all periods as the *low type*, and the other type as the *middle type*.

First, failing to account for selection will bias one’s estimates. The estimated heterogeneity implies very strong selection effects (see Appendix Table A9). The high types are selected out very quickly—under either regime, roughly 90% respond by the first deadline (by the end of period 2), and almost 99% respond by the second deadline (by the end of period 4). In contrast, the low types respond very slowly—under either regime, even after period 5, more than 65% of the low types have not yet responded. Given these strong selection effects, we even find in Table 6 instances in which the aggregate hazard rate decreases from one period to the next in spite of an *increase* in each of the type-specific hazard rates. For example, in the three-type model under the OLD regime, all three types have $p_4^k(\text{OLD}) > p_3^k(\text{OLD})$, reflecting the impact of receiving the first reminder letter. If, however, one considered only a one-type model—thus failing to account for heterogeneity and the resulting selection—one would estimate an average “treatment” effect of $p_4(\text{OLD}) < p_3(\text{OLD})$, which is the opposite direction. Of course, this difference is driven by the significant increase in the proportion of the remaining population that is low type as we move from period 3 to period 4.

Second, our estimates highlight the importance of incidence analysis. Specifically, the different types react to deadlines in different ways. Within each regime, the extent to which people respond in period 2 (just before deadline 1) rather than period 3 (just after deadline 1) reflects the impact of the deadline on their behavior. For the high type, the average daily hazard rate is significantly higher in period 2 than in period 3 (7.47% relative to 5.00% in the OLD regime, 10.43% relative to 9.85% in the NEW regime). For the middle type, the average daily hazard rate actually gets higher in period 3 (3.10% vs. 3.29% in the OLD regime, 4.01% vs. 4.98% in the NEW regime). For the low type, the average daily hazard rate is roughly the same in periods 2 and 3 under either regime. Applying a similar approach to analyze behavior before and after deadline 2 (comparing periods 4 and 5), the high types again react strongly, the middle types now have a mild reaction, while the low types react very little. Hence, it seems that the high types react strongly to deadlines, while the middle

and low types react much less.

The different types also react to letters in different ways, where reminders seem to have their primary impact on helping the better types to pay in a very timely way instead of in a mostly timely way. Table 6 reveals that adding NEW letter 1 (relative to the OLD regime) has a dramatic impact on average daily hazard rates for the high types in periods 2 and 3. Average daily hazard rates increase from 7.47% to 10.43% in period 2, and from 5.00% to 9.85% in period 3. NEW letter 1 also has a large impact on average daily hazard rates for the middle types in periods 2 and 3, increasing them from 3.10% to 4.01% in period 2, and from 3.29% to 4.98% in period 3. In contrast, NEW letter 1 has little impact on the low types—average daily hazard rates in periods 2 and 3 increase only marginally, with little impact on the cumulative response in periods 2 and 3 combined. There is perhaps a more noticeable response among the low types to OLD letter 1 in the OLD regime (i.e., comparing behavior in periods 4 and 5 across regimes), although average daily hazard rates are still very low in both regimes—indeed, by the end of period 5, the cumulative response rate of the low types is only 34.9% in the OLD regime and 31.2% in the NEW regime. Hence, our analysis suggests that if one wants to induce higher response rates and fewer late penalties (and especially if one wants to reduce default judgments and selling debts to collection agencies), the focus should be less on deadlines and reminders, and more on identifying the low types and tackling the difficult question of how to get them to act.

5.3 Analysis of Predicted Types

Once we have estimated the model, we are able to generate a predicted type for each of the 657,890 plates used in the estimation. We can then use these predicted types to investigate (i) the daily hazard rates for each type under the OLD and NEW regimes using the holdout sample of 657,890 tickets, one for each plate, (ii) the daily hazard rates for each type under the eight experimental treatments for plates that also receive tickets under the EXP regime, and (iii) whether the predicted types are correlated with observables. Throughout this section, we apply the estimated three-type model (i.e., the $K = 3$ columns in Table 6).

Using the estimated parameters for the π_k 's and the $p_t^k(\gamma)$'s, the predicted probability

that plate i with observed behavior d^i is type k is

$$\hat{\pi}(k|d^i) = \frac{\pi_k \ell_k(d^i)}{\sum_{k'} \pi_{k'} \ell_{k'}(d^i)}.$$

In principle, we could merely assign plate i to the type k that maximizes $\hat{\pi}(k|d^i)$. With this approach, however, one might worry about plates that are barely assigned to one type relative to another. Instead, we assign plate i to the type k that maximizes $\hat{\pi}(k|d^i)$ only if that k yields $\hat{\pi}(k|d^i) > 0.60$. With this approach, we type 582,063 of the 657,890 plates (88.5%). Of these, 36.2% are assigned to be high types, 39.4% to be middle types, and 24.4% to be low types.²²

Using the holdout sample, Figure 6 depicts the type-specific daily hazard rates in the OLD vs. NEW regimes. Comparing Figures 5 and 6, notice that, not surprisingly, our mixture model provides an improved typing relative to the crude approach to typing used for Figure 5—in the sense that the three types appear more distinct in Figure 6. In addition, Figure 6 yields much the same message as the estimated parameters in Table 6.²³ The high and middle types behave qualitatively the same, with the high types being more prone to act sooner, and both types reacting strongly to the reminder letters. The low types, in contrast, exhibit very low and relatively flat response rates from day 0 through the third deadline. There is a barely noticeable reaction to NEW letter 1 in the NEW regime, and a still small but more noticeable reaction to OLD letter 1 in the OLD regime. In fact, the shift from the OLD to the NEW regime clearly leads to worse cumulative outcomes for the low types. Finally, note that the most significant jump in hazard rates for the low types occurs at roughly day 110, just after they receive letter 3—indeed, under each regime, hazard rates for the low types are higher from day 110 to day 135 than they are for any earlier set of dates. We return to this point in Section 6.

²²The criterion $\hat{\pi}(k|d^i) > 0.60$ is chosen to balance sufficient confidence in the typing against typing sufficiently many plates. If we instead require $\hat{\pi}(k|d^i) > 0.50$, we can type 99.3% of plates, whereas if we require $\hat{\pi}(k|d^i) > 0.70$, we can type 72.8% of plates. See Appendix 9.1.

²³We remind the reader that the summonses used to produce Figure 6 were not used to estimate the parameters in Table 6. For more detail, see Figure A2 in Appendix 9.1, which presents Figure 6 with a separate panel for each type.

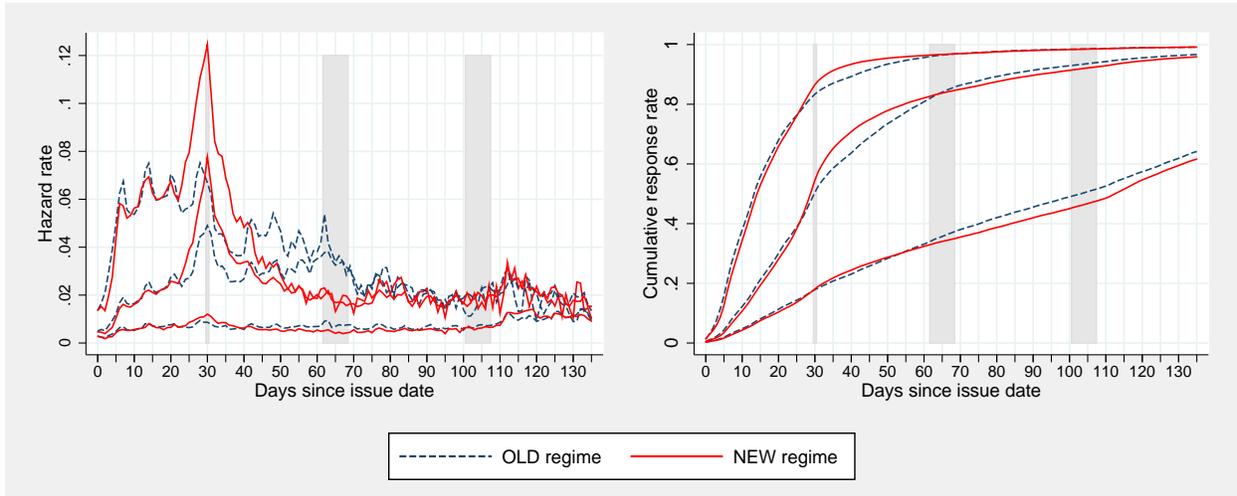


Figure 6: **Hazard Rates and Cumulative Response Rates by Predicted Type.** Note: Daily hazard rates and cumulative response rates in the OLD vs. NEW regimes for each of the three predicted types (high, middle, low) from the estimated three-type mixture model. Based on 582,063 summonses in the holdout sample.

Of the 582,063 typed plates, 78,136 also receive at least one ticket under the EXP regime, and those 78,136 plates receive a total of 98,752 tickets under the EXP regime. Using these tickets, we estimate daily hazard rates in each of the eight experimental cells (see Appendix 9.2 Figures A3 and A4). Due to the much smaller samples, which further shrink as the days pass (especially rapidly for the higher types), our results here are less precise. That said, for each type, the four versions of NEW letter 1 appear to have roughly the same impact on hazard rates, and the EXP letter 1.5 also appears to have a modest impact on hazard rates.

Finally, we assess whether the predicted types are correlated with observables (see Appendix 9.3 for full details, here we provide an overview of our approach and our results). We take advantage of the fact that we have an address for every ticket in the EXP regime for which a NEW letter 1 was sent, and thus we can match the associated plates to Census demographics at the block group level.²⁴ Specifically, of the 657,890 plates in our estimation sample, 60,529 plates received a ticket under the EXP regime for which (i) they were sent a NEW letter 1, (ii) we were able to match their address, and (iii) there were no missing values for the demographic variables.²⁵

²⁴Census block groups have a population of 600 to 3,000 people.

²⁵Whenever a NEW letter 1 is sent to a P.O. box, we treat the address as a non-match. In addition, a small fraction of the addresses were not recognized by ArcGIS, and we treat those as a non-match as well.

For these 60,529 plates, we further update their predicted type probabilities (i.e., their $\hat{\pi}(k|d^i)$'s) based on their response behavior under the EXP regime—this updating is necessary to correct for selection due to the fact that we observe an address only if a ticket is not paid within the first 18 days. Then, in two separate regressions, we regress the predicted likelihood of being a high type and the predicted likelihood of being a low type on various Census variables, including income, race, education, ability to speak English, and how one travels to work (see Appendix 9.3 Table A10 for descriptive statistics and Table A11 for regression results). We find that the probability that a plate is a high type rather than a low type is associated with living in a Census block group that has higher income, more non-native English speakers, a higher proportion of Asian or Caucasian residents, and a higher proportion of college graduates.

6 Discussion

In this paper, we have provided a detailed analysis of response behavior by New York City parking-ticket recipients, estimating daily hazard rates for 135 days after ticket receipt. We have documented clear evidence of forgetting, a central role for heterogeneous types and selection, and key differences in how different types respond to deadlines and reminders. Nonetheless, there is still much to be learned in this domain. We conclude by discussing some issues raised by our analysis, and some thoughts for the future.

First, note that, even after we split our population into types, the type-specific hazard rates in Figure 6 do not look like those in Figure 2. In other words, even after controlling for heterogeneity, we are still left with type-specific hazard rates that are inconsistent with the optimal delay model. Two natural explanations are (i) there is residual heterogeneity within the three types, and (ii) each type is characterized by forgetting, which, as we have emphasized, can generate a general decline in hazard rates.

Second, there are surely other forces at play in this context that we have not studied. One potential force that we have not discussed is another form of inattention: while people have a ticket “on the mind,” they have trouble tracking the exact deadlines that they face. For instance, one might have the ticket on the mind and be looking for a convenient time

to pay, but somehow get the deadline mixed up. A related possibility is that one might open one’s mail at irregular intervals, and thus might not see a reminder letter until after a deadline has passed. Indeed, we see some evidence of this kind of inattention in that response rates remain high in the days after each deadline (i.e., in periods 3 and 5).

Third, our analysis raises the questions of who are the low types and how do we get them to respond. Figure 5 suggests that, at least for some domains, it may be relatively easy to identify the types prone to delay. Hence, with further investigation and more data, it might be possible to learn who they are and why they delay. Perhaps they just do not care about monetary penalties and thus do not care about responding until the consequences become severe. Or perhaps they are liquidity constrained or face other creditors, and thus have limited scope to respond in a timely manner. A more detailed understanding of the low types would then permit programs targeting those types. On this dimension, Figure 5 and especially Figure 6 suggest that some low types react to the third letter—in Figure 6, hazard rates for low types are highest from day 110 through day 135, perhaps due to the more severe consequences or the stated deadline being “IMMEDIATELY.” Hence, it is possible that some types of noticing or language could affect the behavior of low-types. But even these higher hazard rates are still low compared to those for the high or middle types. Hence, other approaches are needed.

Finally, we conclude with a word of caution. A number of recent papers (many cited in Section 1) have identified a significant impact of reminders. We view these results as important, and indeed we find such an impact in our data. We worry, however, that researchers might become overly focused on such unobtrusive interventions to the detriment of investigating when a heavier hand is necessary. We also call for more incidence analysis before giving policy advice. In our domain, the impact of reminders is concentrated on those least likely to need them (i.e., those who would have responded within a few weeks anyway). For the group that delays the most (i.e., the low types, many of whom never respond in our data), our analysis reveals that bolder approaches may be needed.

References

- AKERLOF, G. A. (1991): “Procrastination and Obedience,” *American Economic Review*, 81(2), 1–19.
- ARIELY, D., AND K. WERTENBROCH (2002): “Procrastination, Deadlines, and Performance: Self-Control by Precommitment,” *Psychological Science*, 13(3), 219–224.
- BERTSEKAS, D. (2005): *Dynamic Programming and Optimal Control*. Athena Scientific.
- ERICSON, K. M. (2011): “Forgetting We Forget: Overconfidence and Memory,” *Journal of the European Economic Association*, 9(1), 43–60.
- (forthcoming): “On the Interaction of Memory and Procrastination: Implications for Reminders, Deadlines, and Empirical Estimation,” *Journal of the European Economic Association*.
- HSIAW, A. (2013): “Goal-Setting and Self-Control,” *Journal of Economic Theory*, 148, 601–626.
- KOCH, A. K., AND J. NAFZIGER (2011): “Self-regulation through Goal Setting,” *Management Science*, 113(1), 212–227.
- LAIBSON, D. (1997): “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 112(2), 443–477.
- MCCALL, J. J. (1970): “Economics of Information and Job Search,” *Quarterly Journal of Economics*, 84(1), 113–126.
- MILKMAN, K. L., J. BESHEARS, J. J. CHOI, D. LAIBSON, AND B. C. MADRIAN (2011): “Using implementation intentions prompts to enhance influenza vaccination rates,” *Proceedings of the National Academy of Sciences*, 108(26), 10415–10420.
- NICKERSON, D. W., AND T. ROGERS (2010): “Do You Have a Voting Plan? Implementation Intentions, Voter Turnout, and Organic Plan Making,” *Psychological Science*, 21(2), 194–199.

- O'DONOGHUE, T., AND M. RABIN (1999): "Doing It Now or Later," *American Economic Review*, 89(1), 103–124.
- (2001): "Choice and Procrastination," *Quarterly Journal of Economics*, 116(1), 121–160.
- SALANT, S. W. (1977): "Search Theory and Duration Data: A Theory of Sorts," *Quarterly Journal of Economics*, 91(1), 39–57.
- SCHOAR, A., AND X. CADENA (2011): "Remembering to Pay? Reminders vs. Financial Incentives for Loan Payments," Unpublished, Massachusetts Institute of Technology.
- TASOFF, J., AND R. LETZLER (2014): "Everyone Believes in Redemption: Overoptimism and Nudges," *Journal of Economic Behavior and Organization*, 107A, 107–122.
- TAUBINSKY, D. (2014): "From Intentions to Actions: A Model and Experimental Evidence of Inattentive Choice," Unpublished, Harvard University.
- ZAMIR, E., D. LEWINSOHN-ZAMIR, AND I. RITOV (2015): "It's Now or Never! Using Deadlines as Nudges," Hebrew University of Jerusalem Legal Studies Research Paper Series 16-03.

Table 1: Descriptive Statistics

	Core Dataset	OLD Regime	NEW Regime	EXP Regime
Total # of Tickets	6,646,540	3,355,094	3,020,357	271,089
<u>Violation Type</u>				
Expired Meter (\$35/\$65)	36.23%	37.52%	34.92%	34.88%
Street Cleaning (\$65/\$45)	26.18%	25.38%	27.01%	26.88%
General No Parking Zone (\$65/\$60)	9.21%	9.27%	9.14%	9.28%
General No Standing Zone (\$115)	6.70%	6.58%	6.78%	7.24%
Fire Hydrant (\$115)	5.59%	5.24%	5.95%	5.78%
Double Parking (\$115)	4.75%	4.91%	4.63%	4.00%
Bus Stop (\$115)	2.40%	2.30%	2.50%	2.39%
Truck Loading/Unloading (\$95)	2.17%	2.09%	2.24%	2.22%
Authorized Vehicles Only (\$95/\$65/\$60)	1.94%	2.04%	1.85%	1.73%
In Commercial Zone (\$115)	1.35%	1.25%	1.40%	2.09%
In Crosswalk (\$115)	1.02%	0.90%	1.15%	1.10%
On Sidewalk (\$115)	0.68%	0.70%	0.66%	0.64%
Parking Longer than Limit (\$65/\$60)	0.37%	0.43%	0.32%	0.23%
In a Driveway (\$95)	0.30%	0.30%	0.30%	0.30%
Not as Marked (\$65)	0.23%	0.23%	0.22%	0.30%
In Pedestrian Ramp (\$165)	0.22%	0.20%	0.25%	0.27%
In a Safety Zone (\$115)	0.22%	0.20%	0.24%	0.23%
In a Bike Lane (\$115)	0.17%	0.16%	0.18%	0.17%
No Standing / Taxi Stand (\$115)	0.14%	0.13%	0.15%	0.16%
In Handicapped Zone (\$180)	0.13%	0.16%	0.11%	0.11%
<u>Ticket Amount</u>				
\$35	30.11%	31.30%	28.91%	28.70%
\$45	23.89%	23.12%	24.69%	24.47%
\$60	8.20%	8.25%	8.15%	8.29%
\$65	10.45%	10.64%	10.23%	10.49%
\$95	3.97%	3.95%	4.00%	3.87%
\$115	23.00%	22.36%	23.64%	23.80%
\$165	0.22%	0.20%	0.25%	0.27%
\$180	0.13%	0.16%	0.11%	0.11%
Other/Missing	0.02%	0.02%	0.01%	0.01%
<u>Ticket Issuer</u>				
Parking-Ticket Agent	97.16%	97.28%	97.03%	96.98%
New York City Police Department	2.84%	2.72%	2.97%	3.02%
<u>Payment Type</u>				
Payment made by Day 135	5,333,147	2,721,947	2,397,666	213,534
Mail	32.34%	33.50%	31.23%	29.94%
Online	53.81%	51.11%	56.55%	57.48%
Phone	2.76%	2.10%	3.36%	4.33%
In Person	11.09%	13.28%	8.85%	8.25%
Unknown	0.00%	0.00%	0.00%	0.00%

Note: For all but payment type, percentages in each column are relative to the total number of tickets for that regime (listed on line 1). For payment type, percentages in each column are relative to number of tickets with payment made by Day 135.

Table 2: Timeline For Each Regime

Day	Event	OLD Regime	NEW Regime	EXP Regime
0	Ticket received	✓	✓	✓
19-21	NEW letter 1 (sent on Day 19 unless weekend)	--	✓	✓ (4 versions)
30	Deadline 1 (\$10 late penalty)	✓	✓	✓
35-41	OLD letter 1 (sent on Tuesday)	✓	--	--
47-49	EXP letter 1.5 (sent on Day 47 unless weekend)	--	--	✓ (to 50%)
62-68	Deadline 2 (on Monday, \$20 late penalty)	✓	✓	✓
70-76	Letter 2 (sent on Tuesday)	✓	✓	✓
101-107	Deadline 3 (on Friday, \$30 late penalty)	✓	✓	✓
105-111	Letter 3 (sent on Tuesday)	✓	✓	✓

Note: The shaded areas indicate the communications corresponding to the regime-shift and experimental variations. See details in text.

Table 3: Summary Statistics for First Responses

	OLD Regime		OLD-post Regime		NEW Regime		EXP Regime	
Total # of Tickets	3,355,094		1,240,286		3,020,357		271,089	
Payments by Day 135	61.7%		69.9%		70.8%		70.3%	
Contests by Day 135	28.4%		19.3%		18.5%		19.2%	
Regular Contests by Day 135	16.6%		19.3%		18.5%		19.2%	
Settlements by Day 135	11.8%		0.0%		0.0%		0.0%	
No Response by Day 135	9.9%		10.8%		10.7%		10.5%	
	% of	% of	% of	% of	% of	% of	% of	% of
	Total	Interval	Total	Interval	Total	Interval	Total	Interval
Prior to Deadline 1	53.1%	-----	52.7%	-----	56.2%	-----	55.9%	-----
Payments	36.0%	67.7%	41.1%	77.9%	44.6%	79.4%	43.8%	78.4%
Contests	17.1%	32.3%	11.7%	22.1%	11.6%	20.6%	12.1%	21.6%
Between Deadlines 1 & 2	22.9%	-----	22.6%	-----	19.6%	-----	20.3%	-----
Payments	16.1%	70.2%	18.3%	80.8%	16.0%	81.5%	16.5%	81.5%
Contests	6.8%	29.8%	4.4%	19.2%	3.6%	18.5%	3.7%	18.5%
Between Deadlines 2 & 3	8.8%	-----	8.4%	-----	7.7%	-----	7.6%	-----
Payments	6.1%	69.6%	6.8%	80.9%	6.4%	83.5%	6.5%	85.6%
Contests	2.7%	30.4%	1.6%	19.1%	1.3%	16.5%	1.1%	14.4%
Between Deadline 3 & Day 135	5.3%	-----	5.5%	-----	5.8%	-----	5.8%	-----
Payments	3.5%	65.6%	3.8%	69.2%	3.8%	64.7%	3.5%	60.8%
Contests	1.8%	34.4%	1.7%	30.8%	2.1%	35.3%	2.3%	39.2%

Note: In the bottom panel, "% of Total" is the percentage in reference to the number of tickets in the respective regime (column), and "% of Interval" is the percentage in reference to the number of first responses in the respective regime (column) and time interval (row).

Table 4: Responses Analyzed by Period

Definition of Periods (same for both regimes)					
Period	1	2	3	4	5
Start	day 0	day 20-22	31	day 36-42	day 63-69
End	day 19-21	30	day 35-41	day 62-68	day 70-76
Average Daily Hazard Rates					
Period	1	2	3	4	5
OLD	2.28%	2.69%	2.00%	1.86%	1.32%
NEW	2.17%	3.51%	2.88%	1.33%	0.90%
Cumulative Response Rates					
Period	1	2	3	4	5
OLD	37.63%	53.14%	60.17%	76.02%	78.44%
NEW	36.18%	56.18%	65.27%	75.79%	77.48%

Note: See Appendix 6 for details of how average daily hazard rates and cumulative response rates are calculated.

Table 5: Letters Sent in the Eight Experimental Cells

		EXP letter 1.5 Treatment	
		<i>not sent (50%)</i>	<i>sent (50%)</i>
NEW letter 1 Treatment	<i>baseline (20%)</i>	NEW letter 1	NEW letter 1, EXP letter 1.5
	<i>info (40%)</i>	NEW letter 1 <i>i</i>	NEW letter 1 <i>i</i> , EXP letter 1.5
	<i>scary (20%)</i>	NEW letter 1 <i>s</i>	NEW letter 1 <i>s</i> , EXP letter 1.5
	<i>info scary (20%)</i>	NEW letter 1 <i>is</i>	NEW letter 1 <i>is</i> , EXP letter 1.5

Table 6: Estimated Mixture Model with Average Daily Hazard Rates

	OLD K=1	NEW	OLD K=2	NEW	OLD K=3	NEW	OLD K=4	NEW
π_1	1.000 ----		0.640 (0.001)		0.338 (0.001)		0.261 (0.001)	
p ₁	2.26% (0.00%)	2.15% (0.00%)	3.70% (0.01%)	3.45% (0.01%)	6.58% (0.02%)	6.09% (0.02%)	7.99% (0.03%)	7.46% (0.03%)
p ₂	2.43% (0.01%)	3.18% (0.01%)	4.80% (0.01%)	6.46% (0.02%)	7.47% (0.04%)	10.43% (0.06%)	6.46% (0.06%)	9.74% (0.07%)
p ₃	1.99% (0.01%)	2.84% (0.01%)	4.42% (0.02%)	7.30% (0.03%)	5.00% (0.06%)	9.85% (0.11%)	3.33% (0.06%)	7.68% (0.12%)
p ₄	1.84% (0.00%)	1.31% (0.00%)	5.25% (0.02%)	3.85% (0.02%)	5.69% (0.06%)	3.52% (0.06%)	4.99% (0.07%)	2.98% (0.06%)
p ₅	1.30% (0.01%)	0.87% (0.01%)	4.32% (0.05%)	2.60% (0.04%)	2.45% (0.11%)	1.65% (0.10%)	2.40% (0.12%)	1.64% (0.10%)
π_2			0.360 (0.001)		0.413 (0.001)		0.275 (0.002)	
p ₁			0.60% (0.00%)	0.54% (0.00%)	1.50% (0.01%)	1.35% (0.01%)	2.19% (0.01%)	1.92% (0.01%)
p ₂			0.83% (0.00%)	0.96% (0.01%)	3.10% (0.01%)	4.01% (0.02%)	5.73% (0.03%)	7.36% (0.04%)
p ₃			0.87% (0.01%)	1.16% (0.01%)	3.29% (0.02%)	4.98% (0.02%)	6.21% (0.05%)	9.72% (0.08%)
p ₄			1.06% (0.00%)	0.83% (0.00%)	3.90% (0.02%)	2.93% (0.01%)	6.37% (0.06%)	4.71% (0.06%)
p ₅			0.98% (0.01%)	0.67% (0.01%)	3.57% (0.03%)	2.18% (0.02%)	4.62% (0.13%)	2.62% (0.10%)
π_3					0.249 (0.001)		0.296 (0.001)	
p ₁					0.53% (0.00%)	0.47% (0.00%)	1.17% (0.01%)	1.08% (0.01%)
p ₂					0.52% (0.00%)	0.56% (0.01%)	1.51% (0.01%)	1.91% (0.01%)
p ₃					0.51% (0.01%)	0.60% (0.01%)	1.70% (0.01%)	2.55% (0.02%)
p ₄					0.63% (0.00%)	0.49% (0.00%)	2.54% (0.02%)	1.94% (0.01%)
p ₅					0.65% (0.01%)	0.45% (0.01%)	2.51% (0.02%)	1.59% (0.02%)
π_4							0.169 (0.001)	
p ₁							0.38% (0.00%)	0.33% (0.00%)
p ₂							0.40% (0.01%)	0.40% (0.01%)
p ₃							0.35% (0.01%)	0.36% (0.01%)
p ₄							0.36% (0.00%)	0.27% (0.00%)
p ₅							0.39% (0.01%)	0.27% (0.01%)

Note: Estimated average daily hazard rates by period (p_i 's) for each type, as well as estimated proportions of each type (π_k 's). Standard errors in parenthesis. See Appendix 8.2 for details of the transformation from estimated per-period hazard rates to estimated average daily hazard rates.