

Enforcing Speed Limits Using a Randomization Strategy

Michael Dreyfuss
Jerusalem College of Technology
Havaad Haleumi 21,
Jerusalem, Israel.
dreyfuss@g.jct.ac.il

Irit Nowik
Jerusalem College of Technology
Havaad Haleumi 21,
Jerusalem, Israel
nowik@jct.ac.il

Abstract

Traffic police faces the problem of enforcing speed limits under restricted budget. Implementing high Enforcement Thresholds (ET) will ease the work load on the police but will also intensify the problem of speeding. We model this as a game between two players: The police, which wishes that drivers obey the speed limits, and drivers who wish to speed without getting caught. For the police we construct a strategy in which at each stage the ET is randomized between low and high values. We have established analytically and by simulations that this strategy gradually reduces the ET until it converges to the desired speed limit without increasing the work load along the process. Importantly, this method works even if the strategy is known to the drivers. We study the effect of several factors on the convergence rate of the process. Interestingly, we find that increasing the frequency of randomization is more effective in expediting the process than raising the average amount of fines.

1 Introduction

About 1.25 million people die every year as a result of traffic accidents worldwide and the injuries caused by road traffic accidents are the leading cause of death among young people, aged 15-29 years. Road traffic accidents cost countries approximately 3%-5% of their gross national product.

The ([ETSC, 1999]) has identified three main traffic offences which have a direct connection to road safety and thus ought to be targeted in enforcement strategies: speeding, high blood-alcohol concentration and the lack of use of safety belts. Among these risk factors, speed makes the largest single factor contribution to road accidents. On average, between 40% to 50% of the drivers drive faster than the posted speed limit ([OECD, 2006];[Elvik, 2012];[De Pauw et al., 2004]). An increase in average speed is directly related both to the likelihood of an accident and to the severity of the consequences of the accident.

In 1958 the world's first speed measuring device was introduced and the first speed camera was used to enforce traffic speed limits (Gatso Internet site, [Gatso, 2016]). Since then many countries started using speed enforcement cameras. When a driver speeds beyond the speed limit the camera captures the information of the driver and a notice with a fine or court summon is sent to the driver by mail.

Copyright © by the paper's authors. Copying permitted for private and academic purposes.

In: Yu. G. Evtushenko, M. Yu. Khachay, O. V. Khamisov, Yu. A. Kochetov, V.U. Malkova, M.A. Posypkin (eds.): Proceedings of the OPTIMA-2017 Conference, Petrovac, Montenegro, 02-Oct-2017, published at <http://ceur-ws.org>

In this work, the police needs to choose the Enforcement Threshold (ET) to be implemented at each road. Choosing a low ET may indeed educate drivers to use lower speeds in the future, but will burden the police and court system with heavy work load, whereas choosing a high ET will ease on the system's work load but may intensify the problem of speeding, by leaving such behaviour unpunished.

We view the traffic police and the population of drivers as having contradicting goals; The drivers wish to drive with the *highest* speed without getting caught, and the police wishes that drivers will not speed beyond the speed limit. This, in fact, can be modeled as a game between two players; the traffic police, and the population of drivers. We assume that drivers have beliefs regarding the value of ET and that these beliefs shape their behaviour.

Currently, to avoid heavy work load, the police implements an ET that is substantially higher than the recommended speed limit. This policy ignores the problem of speeding. The goal of this work is to provide the police with a dynamical strategy that gradually reduces the ET until it reaches the minimal value (which is the speed limit itself), while all along, the number of tickets is kept under some predefined limit. This is done by constructing a randomization strategy for the choice of ET. At each stage the police repeatedly randomizes between two low and high values. Indeed, when the result of the randomization is low, many tickets are given, burdening on the police and court system. However, this is balanced out at periods in which the randomization result is high and only few tickets are given out.

It is important to note that for the success of this process, there is no need to keep this method secret. Even if the drivers know about this randomization procedure they cannot utilize this knowledge in their favour as long as the *result* of the randomization is kept secret. This idea of using a randomization strategy in order to confuse and intimidate a potential felon may well be applied to many other features of traffic enforcement, such as changing the location of patrol vehicles strategically etc.

2 The Model

We focus on a specific road and denote the speed limit at that road as b . As mentioned, we denote the Enforcement Threshold that is currently enforced by the police as ET. This means that speeding beyond ET leads to penalty.

The notion of 'population of drivers' in this work refers to drivers which use this road on a regular basis (and could thus go through a learning process). We omit drivers who *always* drive below the speed limit, since these drivers do not need to be educated. Hence we are left with potential traffic offenders, namely drivers who with positive probability speed beyond the speed limit.

Regarding the beliefs of the drivers, it is natural to assume that as long as a driver is not penalized, he just keeps on his old behaviour. Additionally, we assume that a driver changes his belief only if he realizes that it is inconsistent with reality. Hence, even if he *was* penalized, but this event was anticipated by him (as his speed was beyond what he believed to be the ET), then again he does not change his beliefs. Summarizing these two arguments; Beliefs are altered only in case of an unexpected penalty. It follows that beliefs are adjusted always downward and never upward. We view this adjustment process as a 'learning' or 'educating' process. The exact change in behaviour of drivers as a response to a change in their beliefs is most certainly different across countries and cultures, and can probably be estimated by the police via past data. For instance, the Israeli traffic police reports that after being penalized, a driver obeys the law for a while and drives below the speed limit b , and only later on does he become a traffic offender again. Accordingly, we assume that if a driver's speed is v , and he is penalized unexpectedly, then after a period of good behaviour, he adjusts his new belief to v . The relation between beliefs and resulting speeding is presented as a positive density function $f_V(A, v)$, over the speeds $v \in R^+$, where A is the driver's current belief regarding ET, and $V = v$ is his speed. We define, $\bar{F}_V(A, v) = P(V > v|A)$, namely the probability that a random driver will speed beyond v , if he believes that $ET = A$. Hence, $\bar{F}_V(A, v) = \int_v^\infty f_V(A, x)dx$. We make several assumptions regarding $\bar{F}_V(A, v)$. First, we assume that the probability that a driver will speed, enhances when he believes that the enforcing threshold ET , is higher (i.e., more permissive). Hence $\bar{F}_V(A, v)$ is increasing in A , $\forall v$. Secondly, we assume that $\bar{F}_V(A, v)$ is decreasing in v , and $\lim_{v \rightarrow \infty} \bar{F}_V(A, v) = 0$, $\forall A$. This follows immediately from the definition of $\bar{F}_V(A, v)$ as the complementary of the accumulative distribution function of $f_V(A, v)$.

The basic feature (or constraint) of this model is that the work load must be bounded. Thus, we denote by c , the upper bound on the probability to be penalized. This is a predefined parameter limiting the work load on the system. It is determined by the police according to the the volume of traffic at this road, the police and court district limitations, etc.

Note that if the police implements $ET = A$ for a long enough time (as will indeed be dictated by our strategy), then the drivers reveal it, since many times they will speed beyond it, and thus gradually adjust their belief to A . Hence, at this point, $\bar{F}_V(A, A)$ is the probability to be penalized. Now, the basic assumption of this work is that reducing the ET increases the amount of fines and work load on the police and court system. Otherwise, the police can just implement $ET = b$ without giving out an amount of fines that is higher than allowed. This implies that, $\bar{F}_V(A, A)$ is decreasing in A . In particular: $\bar{F}_V(b, b) > c$.

Given an arbitrary road, we summarize the notation and assumptions mentioned here, as follows:

2.1 Notations

- b is the speed limit.
- V = The speed of a randomly chosen driver.
- $f_V(A, v)$ is the density function of the random variable V , when the driver believes that: $ET = A$.
- $\bar{F}_V(A, v) = P(V > v|A)$.
- c = The upper bound on the probability to be penalized.

2.2 Assumptions

1. Each time a driver takes this road, he chooses his speed according to $f_V(A, v)$.
2. $f_V(A, v)$ is known to the police, from accumulative past data.
3. If a driver's speed was v , then in case of unexpected penalty he adjusts his belief to $ET = v$.
4. $\bar{F}_V(A, v)$ is an increasing function of A , $\forall v$.
5. $\bar{F}_V(A, v)$ is decreasing in v , and $\lim_{v \rightarrow \infty} \bar{F}_V(A, v) = 0$, $\forall A$.
6. $\bar{F}_V(b, b) > c$.
7. $\bar{F}_V(A, A)$ is decreasing to zero as $A \rightarrow \infty$.

3 Definition of the Process \mathcal{G}_h

First, choose some $h > 0$, to be the adjustment size of the process.

At stage 1:

1. Find A_1 that satisfies:

$$\bar{F}_V(A_1, A_1) = c. \tag{1}$$

2. Find α_1 , $0 < \alpha_1 < 1$, that satisfies:

$$\alpha_1 \bar{F}_V(A_1, A_1 - h) + \alpha'_1 \bar{F}_V(A_1, A_1 + h) = c, \tag{2}$$

where $\alpha'_1 = 1 - \alpha_1$, and A_1 is the solution for Eq. (1).

3. Now, for the next “learning period”, on a regular basis (e.g., each morning) randomize between $A_1 - h$, and $A_1 + h$, according to the probabilities α_1, α'_1 defined in (2). The result of the randomization will be the ET implemented for the whole period, until a new randomization takes place the following period.
4. Change in beliefs: First, any driver with speed v , s.t., $v > ET$, will receive, by mail, a notice of his fine. If in addition, $v < \beta$, (where β is his current belief regarding ET), then the driver is surprised to get a ticket and thus, he will adjust his belief β to v .

5. We continue the above procedure until the average number of fines goes below c . This gives us the opportunity to lower the ET, and by doing so, the proportion of tickets will grow back to the maximum allowed c . At this point we move to the next stage (see the remark below).

Similarly for all stage $i > 1$: *At stage i :*

6. Given A_{i-1} from the previous stage, find the A_i that solves:

$$\bar{F}_V(A_{i-1} - h, A_i) = c. \quad (3)$$

7. Find the α_i , $0 < \alpha_i < 1$, s.t:

$$\alpha_i \bar{F}_V(A_{i-1} - h, A_i - h) + \alpha'_i \bar{F}_V(A_{i-1} - h, A_i + h) = c, \quad (4)$$

where $\alpha'_i = 1 - \alpha_i$.

8. Repeat stages 3-5 by replacing the index 1 by i .

Remark. To understand the switch we make from the function $\bar{F}_V(A_1, v)$ to the function $\bar{F}_V(A_1 - h, v)$, note that after enough time in which we randomize between $A_1 - h$ and $A_1 + h$, all drivers encounter the low result $ET = A_1 - h$, many times while speeding above it. Consequently, they adjust their belief to their current speed v , which, with time, approaches $A_1 - h$. This still hold at subsequent periods in which $ET = A_1 + h$, since beliefs are never adjusted upward.

Note that for each $h > 0$, \mathcal{G}_h defines a strategy for the police. At each stage i we move to A_i that satisfies: $\bar{F}_V(A_{i-1} - h, A_i) = c$. Thus, if there is an equilibrium e , it satisfies that, $\bar{F}_V(e - h, e) = c$. We wish to find an adjustment size h_b such that $e = b + h_b$ will be the equilibrium point. In that case the process will converge to $e - h_b = b$, since randomizing around $b + h_b$, between the two values b and $b + 2h_b$ ultimately leads the beliefs to the lower outcome which in this case is b .

Note that h_b is the solution for the equation:

$$\bar{F}_V(b, b + h_b) = c.$$

4 Numerical Example

Choose f_V to be:

$$f_V(A, v) = \begin{cases} \gamma_A \frac{g(v)}{G(A) - G(0)}, & \text{if } 0 \leq v < A; \\ (1 - \gamma_A) \lambda e^{-\lambda(v-A)}, & \text{if } v \geq A, \end{cases}$$

where g is the density function of the normal distribution with $\mu = 0.8A$, and $\sigma = 1$, the function G is its cumulative function and: $\gamma_A = 1 - e^{-\mu(A-b)}$. It is easy to verify that f_V satisfies the list of assumptions (see chapter 2.2).

From this we get that for all $v \geq A$:

$$\bar{F}_V(A, v) = e^{(\lambda-\mu)A - \lambda v + \mu b},$$

Hence, by (1) and (3):

$$\begin{aligned} A_1 &= b - (1/\mu) \ln c, \\ A_i &= 1/\lambda(-\ln c + b\mu + (\lambda - \mu)(A_{i-1} - h)), \end{aligned}$$

and by (4):

$$\alpha_i = \frac{ce^{\mu A_i - b\mu} - e^{-\lambda h}}{e^{\lambda h} - e^{-\lambda h}}.$$

Now, h_b solves the equation:

$$\bar{F}_V(b, b + h_b) = c.$$

Hence:

$$h_b = (1/\lambda) \ln(1/c).$$

Choosing: $b = 100, c = 0.3, \lambda = 0.2$, and $\mu = 0.1$, we obtain:

$$\begin{aligned}\bar{F}_V(A, v) &= e^{0.1A - 0.2v + 10}, \\ A_1 &= 100 + 10 \ln(10/3) \sim 112.04, \\ h_{100} &= 5 \ln 10/3 \sim 6.0198,\end{aligned}$$

Note that $b + h_b = 106.03$ is a *unique* equilibrium since:

$$\bar{F}_V(A, A + 5 \ln 10/3) = e^{-0.1A - \ln 10/3 + 10}$$

is *strictly* decreasing and so the equation

$$e^{-0.1A - \ln 10/3 + 10} = 0.3$$

has a unique solution: $A = 100$. Hence the process must converge to the unique equilibrium $100 + h_{100} = 106.03$. Indeed, we get:

$$\{A_i\} = 112.04, 109.03, 107.52, 106.77, 106.4, 106.21, 106.114, 106.067, 106.043, 106.031...$$

And so after a while the drivers believe that:

$$ET \sim \lim_{i \rightarrow \infty} A_i - h_{100} = (100 + h_{100}) - h_{100} = 100,$$

and act accordingly.

5 The Simulations

In this section, we would like to describe the simulations that we have executed and the insights gleaned from the results. The aim of the simulations is to shed light on the factors that influence the duration of the learning process until it converges to b . The simulations can also serve as a basis for the police when trying to implement this learning process in reality. In such case, the police needs to estimate the functions describing the drivers' actual behaviour, and then follow the procedure described in the flowchart presented in Figure 1 ahead. Then it can run the simulations described below and get some insights which will help it to decide on the value of the parameters involved.

For the simulations, an arbitrary intercity road was chosen. The road is one of the entering roads to the high tech city of Tel Aviv, which is equipped with speed cameras, and has a maximum speed limit of 90 km/h. On working days there is an average of 5000 drivers on this road. Due to regulation, the limit for fines is $b = 100$ km/h, which is 10 km/h above the maximum allowed, namely a driver may be penalized only if speeding beyond 100 km/h.

5.1 Simulating the Drivers' Behaviour

Recall that β , is the current belief of the driver regarding the value of ET . Every day, during the 5 working days of the week, each driver "chooses" his speed V as follows: With probability γ_β his speed V satisfies: $V = \beta - V_1$, and with probability $1 - \gamma_\beta$ it satisfies: $V = \beta + V_2$, where V_1 has a Lognormal distribution with expectation of $\mu = 10/(\beta - b)$, and standard deviation of $\sigma = 1$, and V_2 is distributed exponentially with rate $\theta = 5$. The function γ_β is defined as: $\gamma_\beta = 1 - e^{-\lambda(\beta - b)}$, where $b = 100$, and $\lambda = 0.08604$.

If the speed V that was "chosen" by the driver is lower than the current value of ET , then nothing happens, and in the following day, the driver randomizes his speed again according to $f(\beta, v)$. If on the other hand V exceeds ET , then after a while (say, after 10 working days), the driver receives a fine by mail. If $V < \beta$, then this event contradicts his current belief regarding ET , hence he adjusts his belief regarding ET to be: $\beta = V$. According to the police, in the first period after a driver is fined he obeys the law, so we simulate his speed as b . After this period, which we call a period of "good behaviour" (and in short "GB") the effect of the fine wears out to some extent and the driver, although behaving better than before, changes his behaviour and from now on chooses his speed according to $f(\beta, v)$, (with his new belief β). The next day this procedure repeats itself. The process is summarized on the left hand side of the flowchart in Figure 1.

5.2 Simulating the Police's Behaviour

The police determines the actual ET that will be implemented. The goal is to ultimately get ET to be as close as possible to $b + h_b$, at the end of the process, since (as explained), in that case all drivers believe that $ET \sim b$. The police will first need to calculate h_b . It will also need to decide the frequency of adjustment (FR) (say, once every 10 working days). To get started, the police needs to calculate $A = A_1$ and $\alpha = \alpha_1$, as dictated at the first stage of the process \mathcal{G}_h described in Subsection 3. Then, if A 'equals' b , then our goal is achieved and the process is stopped. Technically, by " A equals b " we mean that $A < b + 1.1(h - (\ln c + 0.2h)/\lambda)$. As long as this condition is not satisfied, the police continues as follows: It first randomizes between $A - h$ and $A + h$, according to the first stage of the process described in Subsection 3. The result of the randomization will be the ET implemented until the next randomization takes place after 10 days. After 10 days, before running a new randomization, the police checks if the proportion of fines has already gone below c . If not, it randomizes again around the same A . Otherwise, it adjusts A downward (according to the process), and randomize around the new A . This process is summarized on the right hand side of flowchart in Figure 1.

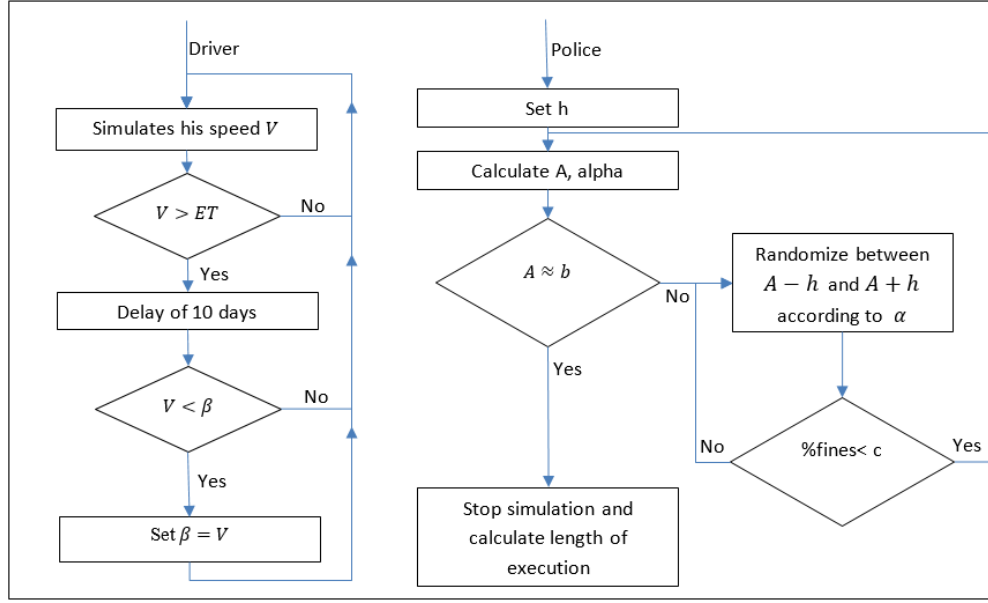


Figure 1: Flowchart of the drivers and police behaviour

The following results focus on the relationships between the execution length (EL), namely the number of days until $A \sim b + h_b$, and the following factors: (1) The duration of good behaviour (GB), (2) The proportion c of fines allowed and (3) the frequency of randomization (FR).

We choose the following base line case:

- It takes two weeks (=10 working days) to transmit the fine to the driver.
- The frequency of randomization is once every 10 working days, namely $FR=10$.
- The upper bound on the percentage of fines is $c=0.2$.
- $GB=40$ working days.

First, let us see how the duration of the simulation (EL) is affected by the duration of good behaviour period (GB). In Figure 2 we see that unsurprisingly, raising GB reduces the EL, i.e, it expedites the learning process. Interestingly, if GB is larger than 10 days, EL is no longer affected by it. A possible explanation for this is that a long enough duration of good behaviour enables us to arrive at the minimal execution time, and this cannot be further improved.




Figure 2: The effect of the duration of good behavior on the duration of the learning process.

In Figure 3 we see the effect of the frequency of randomization (FR) on the duration of the learning process. The smaller the frequency, the longer it takes to educate the drivers. The relationship seems to be linear (for our choice of functions), and we see that this behaviour holds for both cases presented. As shown in Figure 2, at this range, GB has no effect on EL. Hence the difference between the two cases presented in Figure 3 must result from the difference in c .




Figure 3: The relative effect of the amount of fines and the frequency of randomisation on the duration of the learning process.

As expected, the higher c the quicker the process ends. This point is expressed more vividly in Figure 4. In Figure 4 we present the relationship between the amount of fines allowed, and the length of the learning process. Figure 4 shows that raising c allows educating more drivers, thus results in quicker learning process. As seen, this behaviour is the same in both cases. Note that GB in both cases is at least 10 days hence has no effect on EL. Hence the difference between the two cases presented in Figure 4 must result from the difference in FR. As one can expect, higher frequency (FR=10) accelerates the process.

Finally, note that Figure 2 gives an interesting insight on the relative importance of c and FR as factors that influence the duration of the learning process. Recall that raising the frequency of the randomization (i.e., lower FR) accelerates the process. Indeed we see that the solid line, relating to FR=40 is above the dashed line relating to FR=10. However, this happens even though, in the case presented in the solid line, $c=0.5$, which is higher than $c=0.2$ in the dashed line. This means that even though we allow many more fines in the solid line, it still takes longer to educate the drivers (since the frequency of randomization is smaller). The conclusion is that the frequency of randomization is more important than the amount of fines that can be given. Actually these are good news for the police since raising c is expensive, since it leads to more tickets, whereas raising the frequency of randomization can be done automatically, and has no financial consequences.




Figure 4: The effect of the amount of fines on the duration of the learning process.

6 Future Work

From here, there are several interesting directions we wish to take. First, we wish to cooperate with the Israeli traffic police in order to test these ideas on a specific road in Israel. Secondly, we have already formulated the model mathematically, in order to prove that for any reasonable function f , i.e., a function that satisfies our basic conditions, the process is well defined and will always converge to the desired speed limit (when choosing the right adjustment size). Finally, in terms of simulations, it will be interesting to see which of our results still hold for other choices of f . Another interesting direction is to try to apply the ideas underlying our strategy to other features of speed enforcement, e.g., allocation of patrol police cars strategically.

References

- [De Pauw et al., 2004] De Pauw, E., Daniels, S., Brijs, T., Hermans, E. & Wets G. (2014). An evaluation of the traffic safety effect of fixed speed cameras. *Safety Science*, 62, 168-174.
- [Elvik, 2012] Elvik, R. (2012). Speed Limits, Enforcement, and Health Consequences. *Annual Review For Public Health*, 33, 225-238.
- [Elvik, 2015] Elvik, R. (2015). Speed enforcement in Norway: Testing a game-theoretic model of the interaction between drivers and the police. *Accident Analysis and Prevention*, 84, 128-133.
- [ETSC, 1999] ETSC (European Transport Safety Council) (1999). Police enforcement strategies to reduce traffic casualties in Europe, Brussels.
- [Gatso, 2016] Gatso. (2016) <http://www.gatso.com/en/about-gatso/history/> Accessed 23.5.2016.
- [OECD, 2006] OECD. (2006) Organization for Economic Co-operation and Development (OECD). *Speed management*, 55921, Paris: Organization for Economic Co-operation and Development.