Chapter 12
Big Data without Big Brothers: The Potential of Gentle Rule Enforcement

Ido Erev, Manal Hreib, and Kinneret Teodorescu

One of the main goals of laws and regulations is to decrease the frequency of behaviors expected to impair social safety and welfare. These behaviors are defined as violations, and if detected, should be punished. Historically, the main challenge to the design of effective laws and regulations was the difficulty of detecting violations; the low probability of detecting violations undermines the potential benefit to the public good offered by regulatory acts. A common solution to this difficulty involves the use of severe punishments to create deterrence. For example, despite the low probability of actually catching a thief, past enforcers perceived the threat of chopping the thief's hands, or sending them to Australia, as sufficient to reduce thefts.

Becker (1968/2000) shows that under the standard interpretation of rational economic theory, using severe punishments to compensate for insufficient detection should prove highly effective. However, behavioral research has documented deviations from the rational model that challenge the effectiveness of this compensatory approach. One solution to this problem involves the use of advanced big data and surveillance technologies to increase the probability of detection. However, the use of these technologies is often associated with indirect costs in the form of invading privacy. Unwise use of big data for enforcement can give the enforcers too much power and impinge on basic rights.

In the current chapter, we review recent research that sheds light on the costs and benefits associated with the use of big data technologies to enforce laws and rules. In section "The impact of rare events", we summarize basic research on human sensitivity to low-probability (rare) events. We conclude that before gaining experience people are more sensitive to the magnitude of the punishment, but that

I. Erev (✉) · M. Hreib · K. Teodorescu
Faculty of Data and Decision Sciences, Technion Israel Institute of Technology, Haifa, Israel

J. Glückler, R. Panitz (eds.), *Knowledge and Digital Technology*, Knowledge and Space 19, https://doi.org/10.1007/978-3-031-39101-9_12

experience reverses this tendency. The effectiveness of deterrence generated by a threat of severe punishments, therefore, should be short lived. Experienced agents cannot be so easily threatened and are likely to be more sensitive to the probability of detection than to the magnitude of the punishment.

In section "The value of gentle rule enforcement", we highlight the value of gentle rule enforcement. We suggest that severe punishment can be costly for the enforcers themselves, interfering with proper enforcement. Consequently, if the probability of detection can be raised sufficiently, gentle enforcement is more effective than severe punishment. In section "Privacy", we demonstrate that in many settings gentle rule enforcement can be performed with minimal invasion of privacy and does not require changes of current laws. When the probability of the detection of the initial violation is sufficiently high, gentle enforcement can be performed without collecting data about the behavior of specific individuals. In many cases, the focus on the location in space can replace the need to impair privacy. In section "Gentle rule enforcement and the law", we consider the legal implications of our analysis.

The Impact of Rare Events

Experimental studies of human decision-making have revealed contradictory deviations from the prediction of rational economic theory. Kahneman and Tversky (1979) noted that part of the contradictions involves the inconsistent impact of low probability (rare) events. They wrote: "Because people are limited in their ability to comprehend and evaluate extreme probabilities, highly unlikely events are either neglected or overweighted, and the difference between high probability and certainty is either neglected or exaggerated" (Kahneman & Tversky, 1979, p. 283).

The Description-Experience Gap

The effort to clarify the impact of rare events reveals a large difference between initial decisions made purely based on a description of the incentive structure, and subsequent decisions made largely based on past experiences. The top panel of Table 12.1 summarizes Kahneman and Tversky's (1979) study of the impact of rare events on decisions from description. The results reveal high sensitivity to the rare (low probability) outcomes. For example, most participants preferred a "sure loss of 5" over a "1 in 1000 chance to lose 5000." This pattern appears to suggest that if our goal is to reduce the frequency of a specific illegal behavior, rare but severe fines (e.g., a fine of 5000 for 1 in 1000 violations) are likely to be more effective than frequent but low fines with the same expected penalty (e.g., a fine of 5 with certainty).

However, other studies (Barron & Erev, 2003; Hertwig, Barron, Weber, & Erev, 2004; Plonsky & Teodorescu, 2020a) have subsequently revealed that experience

Table 12.1 Comparison of studies of decisions from description with and without feedback

Study	Main results
Decisions from description (Kahneman and Tversky 1979) Method: The participants were asked to choose once between the following two hypothetical prospects: S: Sure loss of 5 R: A 1 in 1000 chance to lose 5000; no loss otherwise	Choice rate of Option S: 80% Under prospect theory (Kahneman and Tversky 1979), this choice rate suggests that most subjects behave as if the probability of the rare event (-5000) is overweighted
The impact of feedback (Erev et al. 2017) Method: In each of 25 trials, the participants were asked to choose once between the following prospects. They were paid (in shekels) for one randomly selected choice, and starting at Trial 6, received full feedback (saw the realized payoffs) after each choice S: Sure loss of 1 R: 1 in 20 chance to lose 20; no loss otherwise	Initial tendency to choose S (52% before receiving feedback), and a reversal of this tendency after several trials. The availability of feedback increased the choice rate of Option R from 48% to 64%

Note. Source: Design by authors

can reverse the impact of rare outcomes. The bottom panel of Table 12.1 presents one demonstration of this observation: When people face repeated choices between a “sure loss of 1” and “1 in 20 chance to lose 20,” they initially tend to prefer the sure loss; after fewer than 5 trials with feedback, however, they change their preference to favor the riskier prospect. Accordingly, the tendency to overweight rare events when considering the initial description is reversed when basing decisions on repeated experiences, leading to under-weighting of rare events in the long run. This pattern is known as the “description-experience gap” (Hertwig & Erev, 2009).

The Reliance on Small Samples Hypothesis and the Intuitive Classifier Explanation

Hertwig et al. (2004) noted that the tendency to underweight rare events in decisions from experience can be captured by assuming that decision-makers rely on only small samples of their past experiences. To see why reliance on small samples will lead to underweighting of rare events, note that the probability that a small sample will not include events that occur with probability $p < 0.5$ tends to be larger than 0.5. Specifically, most samples of size k **will not** include a rare event (that occurs with probability p) when the following inequality holds: P(no rare event included) = $(1-p)^k > .5$. This inequality implies that $k < \log(0.5)/\log(1\text{-}p)$. For example, when $p = 0.05$, $k < 13.51$. That is, when k is 13 or smaller, most samples do not include the rare event (Teodorescu, Amir, & Erev, 2013). Therefore, if people draw small samples from the true payoff distributions and choose the option with the higher sample mean, in most cases they will choose as if they ignore the possibility that the rare event can actually occur.

The hypothesis that people rely on small samples underlies the most successful models in a series of choice prediction competitions (Erev, Ert, & Roth, 2010a, b, 2017; Plonsky et al., 2019) and can be used to explain many judgement and decision-making phenomena (e.g., Erev & Roth, 2014; Erev, Ert, Plonsky, & Roth, 2023; Fiedler, 2000; Fiedler & Juslin, 2006; Kareev, 2000). Plonsky et al. (2015) demonstrate the descriptive value of this hypothesis can be the product of the fact that it is expected both when the decision-makers try to minimize effort, and when they are highly motivated and use sophisticated computations in an attempt to approximate the optimal strategy.

The effort to minimize effort is likely to trigger reliance on small samples when the sampling process is costly, and the benefit from reliance on large samples is relatively low. This effect is particularly clear in studies that focus on search behavior (Hertwig et al., 2004; Wulff, Mergenthaler-Canseco, & Hertwig, 2018; Ackerman, Douven, Elqayam, & Teodorescu, 2020; Teodorescu, Sang, & Todd, 2018).

When people are motivated to maximize expected return, they are also likely to base each choice on small samples if they have reason to believe that the environment is dynamic (e.g., the probability of gain is determined by a Markov chain). In such cases, one can approximate the optimal strategy by relying on a small sample of the most similar past experiences. The thought experiment presented in Figure 12.1 illustrates this assertion.

It is easy to see that in Figure 12.1's example, the intuition (of intelligent decision-makers) is to base the decision in Trial 16 on only three of the 15 past experiences—those that seem most similar to Trial 16. In this example, similarity is determined by the payoff from Top in the preceding three trials: Trials 4, 8, 12 and 16 are similar, because in all of them the payoff in the preceding three trials was "-1, -1, -1." Examining Plonsky et al.'s (2015) results, one can conclude that the underlying cognitive processes are similar to machine learning classification algorithms (like Random Forest, Breiman, 2001) that classify data based on distinct features. In Figure 12.1's thought experiment, intuition uses the feature "the payoff from Top in the last three trials" as a signal to guide the choice in Trial 16.

(a) Task:
In each trial of the current study, you are asked to choose between "Top" and "Bottom", and earn the payoff that appears on the selected key after your choice. The following table summarizes the results of the first 15 trials. What would you select in trial 16?

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Top	-1	-1	-1	+2	-1	-1	-1	+2	-1	-1	-1	+2	-1	-1	-1	
Bottom	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

(b) Implications:
In trial 16, intuition favors "Top" despite the fact that the average payoff from "Top" over all 15 trials is negative (-0.4). This intuition suggests a tendency to respond to a pattern, and implies that only 3 of the 15 trials (Trials 4, 8 and 12) are used to compute the value from "Top" in trial 16.

Fig. 12.1 A thought experiment. Following Plonsky et al., 2015. Source: Design by authors

Under this “intuitive classifier” (Erev & Marx, 2023) explanation, people are likely to consider wide classes of features as signals and use the feature that provides the best classification of the relevant past experiences. One obvious example involves the use of “traffic light color” as a signal to guide driving behaviors. Most drivers use this signal and stop at red lights to avoid accidents and fines. However, when explicit signals such as a red traffic light are absent, people, in an effort to understand their environment, may rely on many other (sometimes irrelevant) signals to sample subsets of past experiences (e.g., Cohen & Teodorescu, 2022; Plonsky & Teodorescu 2020b). Such signals could be their current mood, or the day of the week. Accordingly, by using the intuitive classifier hypothesis one would predict that even highly motivated people are likely to base their decisions on only a small subset of their previous experiences.

The Value of Gentle Rule Enforcement

The reliance on small samples hypothesis suggests that: (1) the deterrence created by a rare but severe punishment will not be effective for most of the population that has already gained some experience in comparable situations; (2) when it is easy to frequently detect violations of laws and regulations, even gentle fines are enough to ensure compliance. For example, if running a red light saves 80 seconds, a frequent fine of 81 seconds should be enough to eliminate this violation, whereas a severe, but rare, 24-hour detention will have little effect in the long run.

In a recent paper, Teodorescu, Plonsky, Ayal, and Barkan (2021) explicitly examined the above predictions in the simple perceptual task described in Figure 12.2. In each study trial, they presented their participants with dots on a divided screen and asked them to report which side contained more dots. Those who reported more dots on one of the sides received a higher reward (10 points vs. 1 point), regardless of the accuracy of their response. Thus, participants could try to increase their earnings by reporting the more profitable side (that with 10 points), even doing so contradicted the evidence. In the first stage, the researchers did not verify the answer, and reporting the more rewarding 10-points side was always beneficial. In the second stage, they informed the participants that from now on, they would randomly sample and verify answers, meting out fines for each incorrect response. As a deterrent, they implemented a policy of high enforcement frequency ($p = 0.9$) with small fines (-10) for one group, and a policy of low enforcement frequency ($p = 0.1$) with high fines (-90) for the other. Notice that the expected value for misreporting was identical in both enforcement policies.

The results revealed that a higher frequency of gentle punishments decreased the rate of violation much more effectively than a lower frequency of more severe punishments. The gap was especially large among particularly delinquent participants (those who tended to commit more violations in the first, non-enforced stage). Moreover, this trend held steady even when the researchers told the participants how much the fine was in advance but did not reveal the frequency of

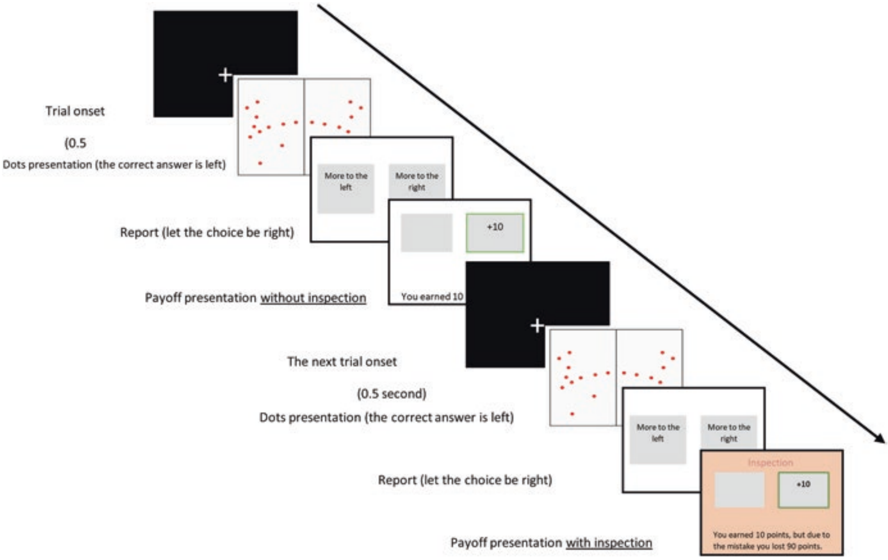

Fig. 12.2 Timeline example of two trials: The first trial without inspection and the second with inspection under an enforcement policy with severe punishment (fine = −90 points). Source: Reprinted from Frequency of enforcement is more important than the severity of punishment in reducing violation behaviors, by Teodorescu et al. (2021, p. 3). Copyright by authors. Reprinted with permission

enforcement—which simulates many real-life situations. From a practical standpoint, one can conclude that when the inspection rate is low, policymakers should prioritize increasing the frequency of inspections over the severity of punishments.

Moreover, as law enforcers are often reluctant to give very large fines, when the expected punishment is severe, law enforcement agents tend to let people go with just a warning. Therefore, large fines could result in a perception of unfairness and consequently reduce the probability of detection (Feess, Schildberg-Hörisch, Schramm, & Wohlschlegel, 2018; Polinsky & Shavell, 2000), which seems to be the key factor in reducing delinquent behavior. Accordingly, these findings are a strong indicator that "gentle rule enforcement" (Erev, Ingram, Raz, & Shany, 2010c) that includes smaller punishments with higher probability would be more effective in reducing violation rates, especially for high offenders, the target population of any enforcement policy.

In order to clarify the significance of this suggestion, it is constructive to note that many substantial violations begin with much lighter breaches. For example, certain cheating efforts, during exams, start with looking around to identify a visible exam form with completed answers. Similarly, certain violent fights in public areas start with carrying concealed weapons, and threatening others with this weapon. The current logic suggests that enforcers can use gentle rule enforcement to stop the first stages in these event sequences that, left untouched, might snowball into a serious violation. In contrast, it is often impossible (or too costly) to stop the first stages with harsh punishments.

In one examination of the value of gentle rule enforcement, Erev et al. (2010c) tried to reduce cheating on college exams. They ran an experiment during the final semester exams of undergraduate courses at the Technion. Traditionally, instructions for exam proctors at the Technion included the following points:

1. The student's ID should be collected at the beginning of the exam.
2. A map of students' seating should be prepared.

As collecting IDs is the first step to constructing this map, proctors commonly interpreted these instructions to mean that they should prepare the map at the start of the exam. Early map preparation was designed to ensure that it will be possible to detect and severely punish cheaters. However, it distracts the proctors and reduces the probability of early gentle punishment (e.g., warning or moving the suspected student to the first row). The experiment compared two conditions that differed with respect to the timing of the map's preparation. In the control condition, the proctors were asked to prepare the map at the beginning of the exam (as they had traditionally done prior to the study), and in the experimental condition, the proctors were asked to delay the preparation by 50 minutes, implicitly allowing them to focus on early detection of cheating intentions. Seven undergraduate courses were selected to participate in the study. In all courses, the final exam was conducted in two rooms. One room was randomly assigned to the experimental and the second to the control condition. After finishing the exam, students were asked to complete a brief questionnaire in which they rated the extent to which students cheated in this exam relative to other exams. The results reveal a large and consistent difference between the two conditions. The perceived level of cheating was lower in the experimental condition in all seven comparisons.

Another examination of the value of gentle enforcement, conducted by Schurr, Rodensky, and Erev (2014), was focused on an attempt to increase compliance with safety rules. Foremen in 11 Israeli factories were asked to encourage the use of safety devices by simply telling workers who did not use them to cease their current work and bring the missing safety devices. This gentle but frequent enforcement mechanism replaced a harsh one in which large fines were occasionally administered by the factories' safety inspectors. The results revealed a quick decrease, from 50% to 10%, in safety rule violations.

To summarize, given people's tendency to rely on small samples of past experiences and the associated sensitivity to enforcement frequency, gentle, yet frequent, rule enforcement seems to be the key to effectively reducing undesired violation behaviors. Although the cost of close monitoring used to be high, recent technological advancements and the increasing usage of AI algorithms enable more effective monitoring with significantly reduced costs (e.g., Abaya, Basa, Sy, Abad, & Dadios, 2014; Piza, Welsh, Farrington, & Thomas, 2019; Raaijmakers, 2019).

Privacy

One of the main risks associated with the use of big data technology for enforcement involves costly invasion of privacy (e.g., Lynch, 2020; Schwartz & Solove, 2011; van Zoonen, 2016). We believe that a gentle rule enforcement policy as discussed above can reduce this risk. Our belief rests on the observation, previously alluded to, that many severe violations start with minor ones. Although identification of people committing severe violations can be important, for minor violations we might prefer to prioritize stopping them early on, without the need to identify the offender. As for most small violations it is not vital to identify the offender, it is thus possible to develop sensors that use big data technology to stop the violation escalating without recording Personalized Identifiable Information (PII). One example of a successful enforcement of this type involves the use of seat-belt alarm systems (Lie, Krafft, Kullgren, & Tingvall, 2008). These systems create an environment where violations of the law "buckle your seat belt" lead to an unpleasant noise with high probability. These systems capitalize on our sensitivity to the frequent event and are thus highly effective despite the fact that they neither collect information about the individuals violating the law nor inflict severe punishments.

Another example involves the use of gentle rule enforcement to reduce cheating in exams, described above. This enforcement was performed without collecting information on the individuals who were asked to move to the first row. The move to the first row was effective because it was enforced liberally but served only as a minor punishment (for example, it wasted time), and also because it served as a frequent, implicit warning.

These examples demonstrate that when the detection probability of the first stage of a sequence of violations is sufficiently high, certain warnings can replace both punishment and invasion of privacy. In order to clarify the potential of this observation, consider the use of video surveillance systems to reduce violence in public areas. Previous research (see Welsh & Farrington, 2009) shows that surveillance systems are rather effective in reducing car related crimes, but much less effective in reducing physically harmful forms of violence (e.g., homicides, fights with injuries, aggravated assaults) in public areas. Under the reliance on small samples hypothesis, this gap in the effectiveness of surveillance cameras reflects the probability of detection (Hreib, 2017). When a car is stolen or damaged, the owner is likely to file a complaint, and the data collected by the surveillance systems significantly increases the probability of identifying and punishing the offender. In contrast, currently, violence is likely to be detected only in the case of serious injuries or homicides. Take, for example, cases in which youngsters use a concealed weapon to threaten others. It is natural to assume that this behavior will usually prove effective: The threatened party is likely to understand the message and back down. In such cases, the existence of the surveillance camera is ineffective because the violation will not be detected.

To illustrate this problem, consider a city with 200 public areas that are covered with video surveillance systems. Assume further that all 200 cameras are connected to an operation room, and two operators monitor the 200 screens with the intention of intervening (sending police) when they detect the beginning of a fight. It is natural to assume (and the reliance on small samples hypothesis would lead one to predict) that the operators are likely to focus on the most interesting screen—the one attached to their smartphone. Thus, the probability of detecting violence in real time is very low. Big data technology can solve this problem. For example, developers can create machine-learning algorithms that detect evidence of threats that include concealed weapons and other indications of the beginning of a fight, and immediately send a warning signal. The signal, say a blue light, can appear both on the screen (in the operation room) and on the camera in the public area. The signal on the screen will draw the operator's attention, and the signal on the camera will inform the fighting parties that the police are on their way. Thus, like the seat belt alarm, it reduces the benefit of violating the law and can stop the violation without collecting Personalized Identifiable Information (PII).

Similarly, undesired smoking in public areas can be detected via smoke sensors, but instead of identifying the individual offender, an automatic reaction can interrupt the smoker. For example, imagine that each time a sensor detects cigarette smoke in a pub, it automatically turns off all lights within a given radius of the detected smoke (or alternatively, turns off the lights by all other tables, leaving light only on the smoking table). In a similar vein, sensors can detect pedestrians running a red light in crosswalks and provide them with an aversive sound (which will also direct nearby people's attention to the violation). More advanced sensors can be used to detect violations such as littering. Imagine that each time something falls from someone's hands, a nearby speaker announces: "Something has fallen on the floor, please pick it back up."

More generally, we suggest that the solutions to many violations start with the use of local sensors to detect the existence of violations in public spaces. Once a sensor has detected a violation, it can send non-private information about its location while simultaneously creating an immediate automatic reaction that signals to the offenders that their violation has been noticed. Thus, by focusing on the space, it can limit the impairment of privacy and direct patrols to where they will be most effective. We suggest that this type of solution generates gentle enforcement, which we expect to reduce small violations (that can lead to serious violations) in public areas without invading privacy.

Gentle Rule Enforcement and the Law

The examples presented above demonstrate that the use of technology to facilitate gentle rule enforcement in public areas does not require new legislation. For example, adding blue warning lights to surveillance cameras does not change the information these cameras collect, nor does it change the punishment meted out to individuals found to violate specific laws. It only directs the attention of the human operator observing multiple screens to a region of interest, consequently increasing the probability of detecting initial violations. At the same time, the blue light at the location itself warns individuals that have begun violating the law that the police are on their way, thus potentially interrupting or preventing more severe violations before they ever occur. We expect these changes to facilitate the enforcement of current laws and regulation in public areas and increase compliance with the law. In addition to reducing severe crimes, we expect them to reduce the necessity of severe punishments.

Yet some violation behaviors occur in private areas (one's car or house), where privacy concerns bar policy-makers from installing sensors linked to automatic responses. In these cases, regulation that forces installation of such sensors in privately owned consumer products can be of help. The most trivial example is the regulation forcing car manufacturers to install sensors that react with an annoying sound when passengers fail to fasten their seat belts (but without reporting this to any central agency). We expect that extending such regulation to additional sensors that detect and react to other dangerous driving behaviors (e.g., driving above the speed limit, changing lanes too frequently, dazzling drivers with strong headlights, etc.) will drastically reduce these violation behaviors. Importantly, in the absence of such regulations, another solution is to incentivize individuals to voluntarily install gentle enforcement devices/apps by, for example, offering discounts on insurance plans to consumers who make use of them.

Summary

Basic decision research suggests that with experience, people become highly sensitive to the most frequent outcomes and tend to underweight rare outcomes. Therefore, rare severe punishments lose their deterrence in the long run. As such, gentle enforcement with high probability is likely to prove more effective in reducing violation behaviors. Big data technologies in surveillance systems and advanced sensors enable substantial increase in the probability of detecting violations, yet they are criticized for invading privacy. The current analysis suggests that these problems can be addressed by building on the observation that most crimes start with small violation behaviors which can be detected and stopped without collecting Personal Identifying Information (PII). Thus, it is possible to develop big data technologies that gently prevent crime and avoid the Big Brother problem.

Acknowledgements This research was supported by two grants from the Israel Science Foundation: Grants 535/17 and 861/22 to Ido Erev, and Grant 2740/20 to Kinneret Teodorescu.

References

Abaya, W. F., Basa, J., Sy, M., Abad, A. C., & Dadios, E. P. (2014). *Low cost smart security camera with night vision capability using Raspberry Pi and OpenCV.* 2014 International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment and Management (HNICEM). Palawan, Philippines: Institute of Electrical and Electronics Engineers. https://doi.org/10.1109/HNICEM.2014.7016253

Ackerman, R., Douven, I., Elqayam, S., & Teodorescu, K. (2020). Satisficing, meta-reasoning, and the rationality of further deliberation. In S. Elqayam, I. Douven, J. St B. T. Evans, & N. Cruz (Eds.), *Logic and uncertainty in the human mind: A tribute to David E. Over* (pp. 10–26). London: Routledge. https://doi.org/10.4324/9781315111902-2

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making, 16*(3), 215–233. https://doi.org/10.1002/bdm.443

Becker, G. S. (2000). Crime and punishment: An economic approach. In N. G. Fielding, A. Clarke, R. Witt (Eds.), *The economic dimensions of crime* (pp. 13–68). London: Palgrave Macmillan. https://doi.org/10.1007/978-1-349-62853-7_2 (Original work published 1968).

Breiman, L. (2001). Random forests. *Machine Learning, 45,* 5–32. https://doi.org/10.1023/A:1010933404324

Cohen, D., & Teodorescu, K. (2022). On the effect of perceived patterns in decisions from sampling. *Decision*, *9*(1), 21–42. https://doi.org/10.1037/dec0000159

Erev, I., Ert, E., & Roth, A. E. (2010a). A choice prediction competition for market entry games: An introduction. *Games, 1*(2), 117–136. https://doi.org/10.3390/g1020117

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., Hertwig, R., Stewart, T., West, R., & Lebiere, C. (2010b). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making, 23*(1), 15-47. https://doi.org/10.1002/bdm.683

Erev, I., Ingram, P., Raz, O., & Shany, D. (2010c). Continuous punishment and the potential of gentle rule enforcement. *Behavioural Processes, 84*(1), 366–371. https://doi.org/10.1016/j.beproc.2010.01.008

Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review, 124*(4), 369–409. https://doi.org/10.1037/rev0000062

Erev, I., Ert, E., Plonsky, O., & Roth, Y. (2023). Contradictory deviations from maximization: Environment-specific biases, or reflections of basic properties of human learning? *Psychological Review, 130*(3), 640–676. https://doi.org/10.1037/rev0000415

Erev, I., & Marx, A. (2023). Humans as intuitive classifiers. *Frontiers in Psychology*, *13*, 1041737.

Erev, I., & Roth, A. E. (2014). Maximization, learning, and economic behavior. *Proceedings of the National Academy of Sciences, 111,* 10818–10825. https://doi.org/10.1073/pnas.1402846111

Feess, E., Schildberg-Hörisch, H., Schramm, M., & Wohlschlegel, A. (2018). The impact of fine size and uncertainty on punishment and deterrence: Theory and evidence from the laboratory. *Journal of Economic Behavior & Organization, 149,* 58–73. https://doi.org/10.1016/j.jebo.2018.02.021

Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review, 107*(4), 659–676. https://doi.org/10.1037/0033-295X.107.4.659

Fiedler, K., & Juslin, P. (Eds.). (2006). *Information sampling and adaptive cognition*. New York, NY: Cambridge University Press. https://doi.org/10.1017/CBO9780511614576

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science, 15*(8), 534–539. https://doi.org/10.1111/j.0956-7976.2004.00715.x

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences, 13*(12), 517–523. https://doi.org/10.1016/j.tics.2009.09.004

Hreib, M. (2017). *Video surveillance and the impact of feedback* (unpublished doctoral dissertation). Technion–Israel Institute of Technology, Faculty of Data and Decision Sciences, Haifa, Israel.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*(2)*,* 263–292. https://doi.org/10.2307/1914185

Kareev, Y. (2000). Seven (indeed, plus or minus two) and the detection of correlations. *Psychological Review, 107*(2)*,* 397–402. https://doi.org/10.1037/0033-295x.107.2.397

Lie, A., Krafft, M., Kullgren, A., & Tingvall, C. (2008). Intelligent seat belt reminders—do they change driver seat belt use in europe? *Traffic Injury Prevention, 9*(5), 446-449. https://doi.org/10.1080/15389580802149690

Lynch, J. (2020). Face off: Law enforcement use of face recognition technology. Retrieved from SSRN. https://doi.org/10.2139/ssrn.3909038

Piza, E. L., Welsh, B. C., Farrington, D. P., & Thomas, A. L. (2019). CCTV surveillance for crime prevention: A 40-year systematic review with meta-analysis. *Criminology & Public Policy, 18*(1)*,* 135–159. https://doi.org/10.1111/1745-9133.12419

Polinsky, A. M., & Shavell, S. (2000). The fairness of sanctions: Some implications for optimal enforcement policy. *American Law and Economics Review, 2*(2)*,* 223–237. https://doi.org/10.1093/aler/2.2.223

Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review, 122*(4)*,* 621–647. https://doi.org/10.1037/a0039413

Plonsky, O., Apel, R., Ert, E., Tennenholtz, M., Bourgin, D., Peterson, J. C., Reichman, D., Griffiths, T. L., Russel, S. J., Carter, E. C., Cavanagh, J. F., & Erev, I. (2019). Predicting human decisions with behavioral theories and machine learning. Retrieved from arXiv. https://doi.org/10.48550/arXiv.1904.06866

Plonsky, O., & Teodorescu, K. (2020a). The influence of biased exposure to forgone outcomes. *Journal of Behavioral Decision Making, 33*(3)*,* 393–407. https://doi.org/10.1002/bdm.2168

Plonsky, O., & Teodorescu, K. (2020b). Perceived patterns in decisions from experience and their influence on choice variability and policy diversification: A response to Ashby, Konstantinidis, & Yechiam, 2017. *Acta Psychologica, 202,* 102953. https://doi.org/10.1016/j.actpsy.2019.102953

Raaijmakers, S. (2019). Artificial intelligence for law enforcement: Challenges and opportunities. *IEEE Security & Privacy*, *17*(5)*,* 74–77. https://doi.org/10.1109/MSEC.2019.2925649

Schurr, A., Rodensky, D., & Erev, I. (2014). The effect of unpleasant experiences on evaluation and behavior. *Journal of Economic Behavior & Organization, 106*, 1–9. https://doi.org/10.1016/j.jebo.2014.05.012

Schwartz, P. M., & Solove, D. J. (2011). The PII problem: Privacy and a new concept of personally identifiable information. *New York University Law Review, 86*, 1814–1894.

Teodorescu, K., Amir, M., & Erev, I. (2013). The experience–description gap and the role of the inter decision interval. *Progress in Brain Research, 202,* 99–115. https://doi.org/10.1016/B978-0-444-62604-2.00006-X

Teodorescu, K., Sang, K., & Todd, P. M. (2018). Post-decision search in repeated and variable environments. *Judgment & Decision Making, 13*(5)*,* 484–500. https://doi.org/10.1017/S1930297500008767

Teodorescu, K., Plonsky, O., Ayal, S., & Barkan, R. (2021). Frequency of enforcement is more important than the severity of punishment in reducing violation behaviors. *Proceedings of the National Academy of Sciences*, *118*(42), e2108507118. https://doi.org/10.1073/pnas.2108507118

van Zoonen, L. (2016). Privacy concerns in smart cities. *Government Information Quarterly, 33*(3)*,* 472–480. https://doi.org/10.1016/j.giq.2016.06.004

Welsh, B. C., & Farrington, D. P. (2009). Public area CCTV and crime prevention: An updated systematic review and meta-analysis. *Justice Quarterly, 26*(4), 716–745. https://doi.org/10.1080/07418820802506206

Wulff, D. U., Mergenthaler-Canseco, M., & Hertwig, R. (2018). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin, 144*(2), 140–176. https://doi.org/10.1037/bul0000115

Ido Erev (PhD in Psychology from UNC Chapel Hill, 1990) is the ATS' Women's Division Professor (2006) in the Faculty of Data and Decision Sciences in the Technion, and an ex-President of the European Association for Decisions Making. His research clarifies the conditions under which wise incentive systems can solve behavioral and social problems.

Manal Hreib (PhD in Behavioral Sciences from Technion, 2017) is the Head of Research, Assessment & Evaluation department at Co-Impact (The Partnership for a Breakthrough in Arab Employment in Israel), and head of Research department at the Negev Coexistence Forum for Civil Equality. Her work mainly focuses on diversity, integration and inclusion of minorities in the Israeli society.

Kinneret Teodorescu is an associate professor at the Technion – Israel Institute of Technology. She received her PhD in Behavioral Sciences from the Technion in 2014 and returned to the Technion as faculty in 2016 after a post-doc at Indiana University. Kinneret studies the effect of incentives on long-term behaviors, focusing on experience-based decisions in varied contexts such as search and exploration, dishonesty, and healthy lifestyle. With the aim of improving social welfare, Kinneret strives to identify interventions/policies that can effectively mitigate behavioral problems.