MAX PLANCK INSTITUTE
FOR RESEARCH ON
COLLECTIVE GOODS

Tim Friehe
Pascal Langenbach
Murat C. Mungan

DOES THE SEVERITY OF SANCTIONS INFLUENCE LEARNING ABOUT ENFORCEMENT POLICY? EXPERIMENTAL EVIDENCE

The Journal of Legal Studies, 2023

Authors: Tim Friehe
Pascal Langenbach
Murat C. Mungan

Title: Does the Severity of Sanctions Influence Learning about Enforcement Policy? Experimental Evidence

Published in: The Journal of Legal Studies
City: Chicago
Editor: The University of Chicago Law School
Year: 2023
Volume: 52(1)
Pages: 83–106
DOI: https://doi.org/10.1086/720642

Does the Severity of Sanctions Influence Learning about Enforcement Policy? Experimental Evidence

Tim Friehe, Pascal Langenbach, and Murat C. Mungan

Abstract

The literature on law enforcement often assumes that the updating of beliefs regarding the probability of detection is a process that is independent from the severity of the sanction. We test this presumption experimentally, using a taking game in which the probability of detection may be either high or low with commonly known probabilities. Individuals gain information about their probability of detection from their experience in the taking game. Some offenders are punished by a severe sanction, while others are sanctioned only mildly, which causes the experience to differ across subjects. Our analysis reveals that the severity of the sanction influences how individuals update their beliefs about the probability of detection, casting doubt on the widely held presumption that the perceived probability of detection and the magnitude of the sanction are separable.

Tim Friehe is a Professor in the Public Economics Group at the University of Marburg. Pascal Langenbach is Senior Research Fellow at the Max Planck Institute for Research on Collective Goods. Murat C. Mungan is a Professor of Law at George Mason University.
We gratefully acknowledge the comments received from a reviewer and Edward R. Morrison. In addition, we are very grateful for the comments received from Christoph Engel, Simon Dato, Yoan Hermstrüwer, Adrian Hillenbrand, Sebastian Kube, Eric Talley, Kathy Zeiler, and participants in seminars at the Max Planck Institute for Research on Collective Goods and at Bonn University. We thank Francesca Loeckx and Marie-Theres Schier for programming z-Tree and Brian Cooper for language help and translations. We are also thankful to Eugenio Verrina and Adrian Hillenbrand, who shared part of their z-Tree code

1. Introduction

Understanding how people form and update beliefs about enforcement policies is central to important theories of crime (see, for example, Akers and Jennings 2016; Sah 1991; Kelling and Wilson 1982) and to the general theory of deterrence (see, for example, Ben-Shahar 1997). For example, in Kelling and Wilson's broken-windows theory (Kelling and Wilson 1982), individuals are more likely to offend in neighborhoods showing decay because they update their beliefs about the probability of an arrest on the basis of this signal. Yet, despite the centrality of beliefs, little is known about how people form and update their beliefs about enforcement policies and what factors influence the process. Survey-based evidence on the perceived probability of detection (hereafter, detection probability) shows not only marked heterogeneity across survey respondents but also notable intraindividual variability over time (see, for example, Bebchuk and Kaplow 1992; Deffin and Ropaul 2019; Lochner 2007). This evidence supports the idea that one's own detection experiences, as well as those of peers, shape how one's beliefs about detection probabilities are formed and updated (see, for example, Anwar and Loughran 2011; Lochner 2007; Sah 1991). When it comes to how the experience of detection should influence beliefs, Bayes's rule represents the benchmark in both the economics and the law-and-economics literatures (see, for example, Ben-Shahar 1997; Daughety and Reinganum 2020; Deffains and Fluet 2020; Hay and Spier 1997; Kaplow 2011; Köszegi 2006; Sanchirico 2012).

This paper studies whether the severity of the sanction (hereafter, sanction severity) impacts how individuals update their beliefs about the detection probability after they have received a signal, knowing that the sanction's magnitude is irrelevant for belief updating according to Bayes's theorem. We therefore evaluate the validity of the presumption that the perceived detection probability and the sanction's magnitude are separable policy instruments—an assumption central to the large literature on optimal law enforcement (see, for example, Garoupa 1997; Polinsky and Shavell 2007).[1] While that literature acknowledges that potential offenders do not know the precise detection probability and thus will dynamically update their beliefs, it is always assumed that their beliefs are independent of the sanction's magnitude (see, for example, Bebchuk and Kaplow 1992; Ben-Shahar 1997; Shavell 2004). The presumption of separability is also at the heart of the long and ongoing discussion in criminology pertaining to the deterrence effects of the certainty versus the severity of punishment (see, for example, Engel and Nagin 2015; Mungan and Klick 2016).

We experimentally test whether learning about the detection probability can be separated from the sanction severity using a simple and transparent experimental design: in part 1, a potential offender chooses whether to take money from a potential victim, knowing that the detection probability can be either high or low according to commonly known prior beliefs. Our treatment variable is sanction severity: offenders, if detected, are subject to either a high or a low sanction. In part 2, against the backdrop of their enforcement experience, subjects state their beliefs about the likelihood that they are in the high-detection-probability condition.

Although sanction severity should be irrelevant to the updating of beliefs, we find that beliefs about the detection probability are influenced by sanction severity. Conditional on the offender being sanctioned, individuals in the high-sanction treatment predict a higher probability of being detected than individuals in the low-sanction treatment. This finding is consistent with the idea that the experience of detection is perceived to be more important or memorable when it is associated with a more severe sanction (hereafter, high sanction). A more vivid detection experience and greater emotional arousal in the high-sanction treatment is predicted to induce subjects to incorporate the detection signal to a greater extent when updating their belief about being in the high-detection-probability condition.

Our key result strongly suggests that perceptions about detection probabilities are interwoven with the sanction severity. Hence, our analysis rejects the long-maintained hypothesis that sanctions and enforcement efforts are strictly separable, and thus it has important implications for deterrence theory as well as policy analysis.

Our analysis suggests that increasing the sanction's magnitude has a double deterrent effect because it increases the sanction as well as the perceived detection probability.[2] In regulating behavior associated with small expected instances of harm, enforcers often use small sanctions, and warnings might be seen as the extreme case of no sanctions. Some criminal procedures—for example, diversion programs and graduated punishments—can cause first-time offenders to avoid significant punishment for their offenses. The leniency of sanctions in these cases is supported by considerations such as proportionality of punishment and the desire to preserve a meaningful second chance for first-time offenders; this, however, is in conflict with optimal deterrence. Our results might be read as providing arguments in favor of increasing sanctions in these cases, because they imply that the marginal

[1] To be clear, we are aware of papers in the law enforcement literature that show that the probability of conviction and the magnitude of the sanction may be linked. For example, Andreoni (1991) provides an important contribution by highlighting that legal decision makers' greater inhibition in imposing higher sanctions in the face of legal uncertainty may result in a lower expected sanction despite a higher sanction magnitude. We are concerned with the fact that the perception of the detection probability is already influenced by the sanction's magnitude.

[2] In fact, when the sanction also plays an expressive role (see, for example, Cooter 1998), changes in its severity may even have a triple deterrence effect.

benefits of deterrence from increasing sanctions are greater than what one would otherwise assume. However, the double deterrent effect of sanctions also has implications for how a fixed level of deterrence could be reached. Since the sanction severity seems to influence the perceived detection probability, a smaller sanction might be sufficient to reach the target deterrence level, unlike under an analysis based on the presumption that sanctions and detection probabilities are separable instruments.

The structure of the paper is as follows. In Section 2, we discuss the related literature. In Section 3, we present the experimental design and the procedures. In Section 4, we present our results. In Section 5, we provide a discussion of our findings. In Section 6, we conclude with a brief summary of our findings and some policy implications of our research.

2. Literature

We study whether the sanction severity influences how individuals update their beliefs about the applicable detection probability. Therefore, we contribute to the literature that studies how imperfect information about detection probability might affect the behavior of potential offenders and to the literature on the updating of beliefs.

Beliefs are central to economic decision-making. Consequently, the formation and updating of beliefs have attracted a lot of academic attention. Empirical studies document that many individuals update their beliefs differently from Bayesian updating (see, for example, Barron [2021] and the references cited therein). Many analyses document that individuals underrespond to feedback; that is, they show conservatism (see, for example, Benjamin 2019; Buser, Gerhards, and van der Weele 2018). However, the average participant often matches predictions of the Bayesian updating benchmark reasonably well (for example, Barron 2021; El-Gamal and Grether 1995; Gotthard-Real 2017; Holt and Smith 2009).[3]

With respect to systematic departures from the Bayesian benchmark, Tversky and Kahneman (1974) and the subsequent literature established that certain events seem more likely to be remembered and can thus be more influential in belief formation (availability bias).[4] In this context, in terms of the theoretical literature, including Gennaioli and Shleifer (2010) and Mullainathan (2002), Enke, Schwerter, and Zimmermann (2020) provide experimental evidence on the relevance of associative recall for belief formation, showing that cues may induce an overweighting of some pieces of evidence. While in our setting there is no similar role for the selective recall of only a subset of signals or for the substitution often discussed in the context of availability (for example, thinking of unemployment among your friends when asked to report the regional unemployment rate), there may be parallels. We argue that the greater importance or memorability of the detection state when the sanction is high will make it more likely for the enforcement experience to be fully incorporated in the subjects' belief-updating process.

A detection signal makes it more likely that the subject actually is in the high-detection-probability condition. This can be considered bad news when the subject plans future offenses, for instance. Recent experimental work considers the role of good versus bad news for belief formation (for example, Barron 2021; Charness and Dave 2017; Coutts 2019; Ertac 2011; Gotthard-Real 2017; Möbius et al.

[3] Accordingly, contributions from the field of neuroeconomic theory maintain the Bayesian approach in some fashion. For example, Brocas and Carrillo (2012) present a framework based on neurobiology in which individuals simplify information from a signal that can take any value on the unit interval by considering only whether it stems from a category including low signal levels or a category including high signal levels but then correctly apply Bayesian updating.

[4] Lieder, Griffiths, and Hsu (2018), for example, argue that overrepresentation of extreme events in decision-making can be interpreted as a rational use of limited cognitive resources.

2022; Sharot and Garrett 2016). Interestingly, Charness and Dave (2017) report that subjects with a relatively greater stake in one of the two states update more like a Bayesian decision maker, arguing that this results from greater attention to the updating task. These studies examine beliefs when subjects are motivated to prefer one of the states of nature either for ego or for financial reasons. Zimmermann (2020) explores the dynamics of motivated beliefs and uncovers the underlying mechanisms, such as selective recall of information received. In line with this, it could be argued for our design that subjects in the high-sanction condition will believe that there is a higher probability of being in the high-detection-probability condition than subjects in the low-sanction state, because having a high detection probability would be particularly bad news in the high-sanction state (see, for example, Bénabou and Tirole 2016).

Our study also relates to the crime literature dealing with the reality that potential offenders do not know the exact level of the detection probability. A theoretical contribution discusses circumstances under which policy makers may credibly communicate information about the applicable level of the detection probability to potential offenders and finds that such conditions may be very few (Baumann and Friehe 2013). This implies that individuals will remain imperfectly informed about the detection probability in many cases. Sah (1991) assumes that each individual observes the enforcement experience of a set of people to arrive at a belief regarding the detection probability and that this estimate is updated over time. In this respect, Rincke and Traxler (2011) provide empirical evidence that people's beliefs about the detection probability respond to enforcement in their neighborhood, using the context of compliance with TV license fees in Austria. Ben-Shahar (1997) theoretically studies a simple 2-period setup and emphasizes own experience in the (Bayesian) belief-formation process. Anwar and Loughran (2011), Matsueda, Kreager, and Huizinga (2006), and Lochner (2007) provide empirical evidence that generally supports the assumption that individuals on average obey fundamental rules of Bayesian updating and the assumptions made in Sah (1991) and Ben-Shahar (1997). Referring to updating, inter alia, Apel (2013) surveys empirical research concerning the determinants of an individual's perception of the risk of formal sanctions.

Finally, we note that our focus is on probabilities of detection that are nonnegligible, that is, significantly greater than 0. From a normative perspective, the well-known analysis in Becker (1968) implies that the optimal probability of detection may be quite low and may in fact be negligible given very large maximal sanctions. Because our objective is to identify potential impacts that may emerge in many enforcement contexts where the detection probability is nonnegligible (for example, speeding), we refrain from reviewing considerations that cause the Beckerian result to fail (see, for example, Garoupa [1997] or Polinsky and Shavell [2007] for a review). We leave it to future research to study how people update beliefs when the detection probability is close to 0.

3. Experimental Design and Procedures

Our basic design is as follows: A potential offender decides whether to take a fixed amount from a potential victim's endowment (similar to Khadjavi 2015; Rizzolli and Stanca 2012; Schildberg-Hörisch and Strassmair 2012). If the offender takes from the victim, he may subsequently be detected and sanctioned. In our design, players can be in a condition with either a high or a low detection probability. Detection implies punishment, and offenders are subject to either a high sanction or a low sanction. After the taking game and the embedded detection experience, we elicit our main variable of interest: subjects' beliefs about the detection probability.

Table 1 summarizes our design. Subjects were first assigned to the role of either offender or victim, a role they kept throughout the experiment.[5] Part 1 of the experiment consisted of the taking game. In part 2, participants stated their beliefs about the detection probability. Part 3 consisted of another taking game in which offenders were paired with a different victim and subject to the same detection probability as in part 1, whereas the sanction level could again either be high or low. In part 4, subjects completed a postexperimental questionnaire. Participants knew from the start that the experiment comprised four parts and that either part 1, part 2, or part 3 would be payoff relevant in addition to part 4. Instructions for parts 2 and 3 were provided only after subjects had completed part 1. In the following, we describe each part in more detail.

3.1. Part 1: Taking Game 1

Part 1 was subdivided into an online part and a part conducted in the laboratory. Offenders made their taking decisions online, and offenders and victims learned the applicable sanction level online. Whether offenders would be detected and sanctioned was determined in the laboratory. Parts 2–4 took place in the laboratory.[6]

3.1.1. Taking Decision

First, subjects read information about part 1 (as explained in Table 1). They could proceed only after correctly answering control questions. Next, participants learned their role: either offender or victim. Each offender was endowed with 2,400 points, whereas victims received an endowment of 7,200 points. We used these asymmetric endowments to induce a higher taking rate (as in, for example, Khadjavi 2015; Schildberg-Hörisch and Strassmair 2012).

In the next step, offenders had to choose whether to take 2,400 points from the victim (which they kept, as in Polinsky and Shavell [2007], for example, independently of later detection). Offenders knew that the detection probability was either high (p_H = 70 percent) or low (p_L = 20 percent). We denote the scenario in which the high (low) detection probability as scenario which the H (L). The prior probability for scenario H was set to q = 60 percent.

[5] All experimental instructions were neutrally worded: the participants' roles were Person A and Person B, the taking decision was framed as the decision to cause a transfer of points from the other player's account to one's own account, detection was referred to as an audit, and the sanction was framed as a points deduction.

[6] Our motivation for splitting part 1 is explained in more detail below.

Table 1. Experimental Design

Stage	Description
Part 1:	
1	Participants of a matched pair learn their role; the offender (victim) receives an endowment of 2,400 (7,200) points
2	The offender chooses whether to take 2,400 points from the victim
3	The sanction level is drawn from a distribution with a 50 percent chance for the low (high) sanction of 240 (4,800) points and presented to the offender and victim
4	The offender is either detected and sanctioned or not, which is revealed to the offender and the victim; offenders knew that the probability of detection was high (70 percent) or low (20 percent); the prior probability for the high-detection-probability state was 60 percent
Part 2	Participants state their (incentivized) updated belief in the likelihood that the high probability of detection applies to their pair; pairs without enforcement experience of their own receive information about another pair; participants choose from the set {0, 1, . . . , 99, 100}
Part 3:	
1	Participants are matched into new groups; the offender (victim) receives an endowment of 2,400 (7,200) points
2	The offender chooses whether to take 2,400 points from the victim
3	The sanction level is drawn from a distribution with a 50 percent chance of the low (high) sanction of 240 (4,800) points and presented to the offender and the victim
4	The offender is either detected and sanctioned or not, which is revealed to the offender and the victim; the probability of detection from part 1 applies
Part 4	Postexperimental questionnaires are administered (risk, cognitive reflection test, personality, and demographics)

Before the offender's taking decision, all subjects were informed that, on detection, either a high sanction of S_H = 4,800 points or a low sanction of S_L = 240 points would be administered with equal likelihood.[7] Participants learned about the sanction level that was relevant to them only after offenders had made their taking decisions. The information about the sanction level was provided to all participants (that is, all victims and all offenders, independent of their taking decision).

We implemented the subjects' imperfect information about their sanction level to avoid a possible selection effect. Otherwise, the sorting of offenders into the category of taker or nontaker could have depended on the sanction's magnitude. An alternative procedure would have been to conduct one treatment in which, from the start, the sanction was known to be high and another treatment in which the sanction was known to be low. The set of takers in the treatment with the high sanctions would, in all likelihood, differ in many ways from the set of takers in the treatment with the low sanction, and likewise for the sets of nontakers.

One downside of this design choice is greater complexity. Given that the uncertainty about the actual sanction level was relevant only for the choice to take and resolved immediately after that, we do not think that the greater complexity at this stage played a role in the later decisions. However, a second potential downside of our design is that learning about the result of the random draw of the sanction's magnitude might stir up the subjects' emotions. For example, a subject might encounter negative

[7] In parallel to our design, Feess et al. (2018), for example, also use two levels of fines: one below the criminal gain and one above the criminal gain, when studying implications of legal errors.

feelings (of bad luck) when the random draw selects the high instead of the low sanction. Such negative emotions might spill over to the belief elicitation stage. To attenuate the potential effect that emotions stemming purely from the resolution of the uncertainty about the sanction level may have for the updating process, we informed subjects about the applicable sanction severity after the taking decision during the online part of part 1. All participants took the online part at least 1 day before they were invited to the laboratory. To provide better context, we note that Grimm and Mengel (2011) and Neo et al. (2013), for instance, report that a cooling-off period of only 10–15 minutes significantly decreases the likelihood of low offers in the ultimatum game being rejected.[8]
To summarize, all subjects knew the conditions' probabilities and possible detection probabilities from the start. The sanction level was revealed immediately after the offender's taking decision. The subjects' detection probability was not revealed during the experiment.[9] Importantly,

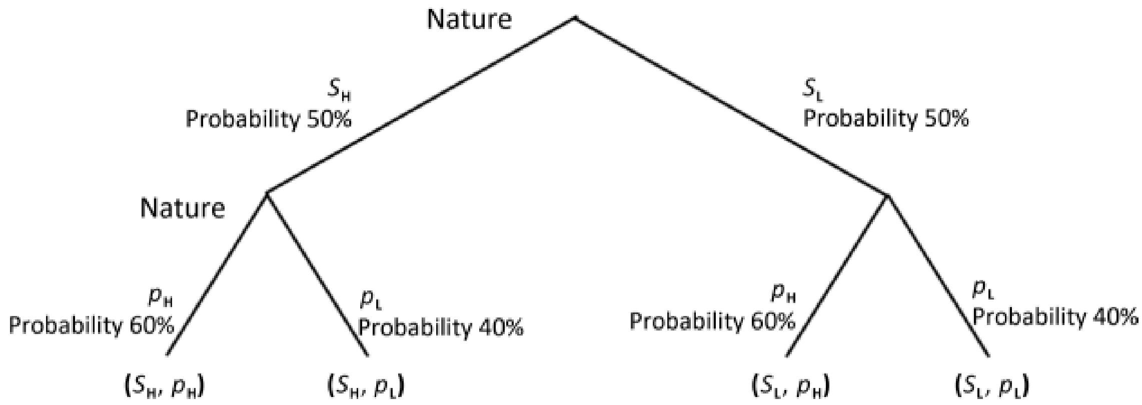

Figure 1. Enforcement policy lotteries in part 1

it should have been clear to subjects that the magnitude of the sanction and the detection probability would be determined in independent draws, leading to one of four possible combinations of sanction and detection probability levels, namely (S_H, p_H), (S_H, p_L), (S_L, p_H), or (S_L, p_L), as shown in Figure 1.

3.1.2. Detection and Sanction

On arrival to the laboratory, subjects received the general information for laboratory experiments, including a recap of the payoff information. To ensure their awareness of the sanction level applicable to them, subjects had to enter the sanction level that had been presented to them during the online part before they could proceed with part 1.[10] Subjects then received an on-screen summary of the decisions made in the online part.
Next, a move of nature determined whether offenders who took from their victims would be sanctioned. In state D (detection), the offender was sanctioned. In state N (no detection), the offender was not detected and thus received 4,800 points from part 1, that is, the 2,400-point endowment plus 2,400 points from the taking. Both the offender and the victim were informed about whether the

[8] Our data provide some indication that there was indeed no direct effect of the resolution of the uncertainty of the sanction on belief updating, since the sanction's magnitude seems to matter only for a subsample in which detection took place (see Section 4) and not per se.

[9] The information about the distributions of the sanction severity and the probability of detection was presented to the subjects in a very transparent way. A translation of the screens that subjects viewed in the online part is in Section OA1 of the Online Appendix.

[10] All subjects succeeded in reproducing the sanction level relevant for them; eight subjects needed more than one trial.

offender was sanctioned. Offenders who decided not to take from the victim were not subject to sanctions; these offenders (and their matched victims) were asked to imagine another offender-victim pair that was in the same detection probability state (either scenario H or L) and the same sanction severity state (either S_H or S_L) but in which the offender took from the victim. For all nontaking pairs, the observed offender-victim pair was in state D. Part 1 ended with information about the offender's detection and sanctioning.

3.2. Part 2: Elicitation of Beliefs

After receiving instructions for parts 2 and 3,[11] the subjects answered control questions concerning the elicitation of beliefs in part 2. The instructions reminded subjects of key aspects of the taking game from part 1, including the levels of the high and low probabilities of detection and the prior probability of having the high probability of detection.

Against this background of enforcement experience, participants stated their beliefs about the likelihood that they were in the group with a high probability of detection (that is, that they were in scenario H in part 1). In part 1, all subjects had been presented with information about the prior probability (that is, q = 60 percent) and had received an informative signal about the probability of detection. Pairs who were involved in a taking in part 1 obtained a signal about the true probability by learning if their offender was sanctioned. Pairs who were not involved in a taking obtained a signal about the true probability by witnessing the enforcement experience of another pair with the same probability of detection and sanction level but in which taking occurred.[12]

Bayesian updating is the standard benchmark and requires subjects to assess how likely it is for them to face the high probability of detection when state N or state D applies to them, starting from the prior probability of 60 percent. For our set of parameters, Bayesian updating yields a posterior probability of 36 percent in state N and 84 percent in state D, which means that a substantial adjustment is due after learning the enforcement outcome.

To incentivize accurate belief reporting by subjects, we closely followed the relevant belief-updating literature and made use of the so-called lottery method.[13] This method has also been used by, among others, Buser, Gerhards, and van der Weele (2018), Coutts (2019), Holt and Smith (2009), and Möbius et al. (2022). The method is designed to ensure that the subjects' optimal response is their true belief, independent of their attitudes toward risk. The lottery method is explained in detail in Section OA3 of the Online Appendix. We made clear that this procedure was implemented to incentivize participants' truthful reporting of their beliefs about the probability of having a high probability of detection. Our instructions and control questions for part 2 closely followed the easily understandable format in Buser, Gerhards, and van der Weele (2018). Participants received information about their earnings from the lottery task only at the end of the experiment.

[11] Instructions for parts 2 and 3 were distributed on paper and read out aloud by the experimenter. We presented experimental instructions for parts 2 and 3 simultaneously to make subjects aware that their probability assessment in part 2 might be relevant later. A translation of the laboratory instructions are in Section OA2 of the Online Appendix.

[12] When no taking occurred, the participants read the following instructions: "Imagine that, in the following, you observe another pair, in which another person A initiated the transfer."

[13] See Schotter and Trevino (2014) and Schlag, Tremewan, and van der Weele (2015) for informative surveys on the elicitation of beliefs.

3.3. Part 3: Second Taking

In part 3, participants maintained their role from part 1 but were matched with a different subject. Offenders chose whether to take 2,400 points from the victim.[14] The probability of detection from the offender's history in part 1 applied, whereas the level of the sanction was again randomly drawn. We implemented part 3 because we wanted to give a practical meaning to the elicitation of probability in part 2 and to explore whether an influence of the sanction level on the updated belief about the probability of detection is behaviorally relevant for future choices.

3.4. Part 4: Questionnaire

In part 4, participants first took an incentivized cognitive reflection test (Frederick 2005) in the expanded version of Toplak, West, and Stanovich (2014). Participants also completed the 10-item Big 5 short inventory (Rammstedt and John 2007) to approximate their personality traits[15] and responded to a validated survey item measuring their risk tolerance (Falk et al. 2016, 2018). Finally, participants provided some sociodemographic information.

3.5. Procedures

The experiment was conducted in the DecisionLab at the Max Planck Institute for Research on Collective Goods in December 2019. Participants were recruited via hroot (Bock, Baetge, and Nicklisch 2014) from the joint subject pool of the DecisionLab and the BonnEconLab of Bonn University (with more than 6,000 subjects). The laboratory part was implemented in z-Tree (Fischbacher 2007).

Two hundred twenty-two subjects participated in 11 sessions in the laboratory.[16] The number of subjects per session varied between 14 and 24 because of varying sign-up and show-up rates and the necessity of matching roles and treatment conditions from the online part. The participants' mean age was 24 years, and 56 percent were female. Almost all of the participants were students (99 percent) from fields such as natural sciences, law, political science, economics, and languages.

Subjects participated in the online part before coming to the laboratory. At least 14 hours (including one night) elapsed between the online and the laboratory parts. Laboratory sessions typically lasted around 80 minutes (including payment). The subjects' earnings were converted to euros using a rate of 350 points to EUR 1 and paid out in cash. Average earnings were around EUR 16.90, including a flat fee of EUR 5 for participation in the experiment.

4. Results

We first briefly describe the offenders' taking behavior in part 1 of the experiment (Section 4.1) and then explain how the stated conditional probabilities in part 2 compare with Bayesian posterior probabilities (Section 4.2.1). Our main research question asks whether severity of the sanction influences people's learning about the probability of detection and is addressed in our analysis in Section 4.2.2. We also briefly look at the offenders' taking behavior in part 3 of the experiment (Section 4.3).[17]

---

[14] Because of the random draw in part 1, part 2, or part 3 as the payoff-relevant part, it is not possible for a victim to experience two takings in terms of payoffs.

[15] The Big 5 personality inventory measures openness, conscientiousness, extraversion, agreeableness, and neuroticism (see, for example, Costa and McCrae 1992). Below we particularly rely on the neuroticism score because this trait relates to emotional stability, which informs us about the extent to which people react to negative stimuli.

[16] Of 303 participants completing the online part, 246 showed up on time. After matching roles and treatment conditions, we collected complete data from 222 participants.

[17] We use the information obtained in part 4 as covariates in the regression exercises.

4.1. Part 1: Taking Decision

Most offenders (76 percent) decided to take points from their victims in part 1. While demographic information is not predictive of taking, the participants' self-reported risk attitude indeed is (ordinary least squares regression, coefficient of .06, p = .001).[18] Therefore, roughly three-quarters of our participants relied on firsthand enforcement experience when updating their beliefs about the probability of detection. The remaining participants used the enforcement experiences of the pair they observed.

4.2. Part 2: Updated Beliefs about the Probability of Detection

Here we first inquire about the match between reported beliefs and Bayesian posterior probabilities. We then address our main research question, that is, whether severity of the sanction influences updated beliefs.

4.2.1. Comparing Stated Conditional Probabilities to Bayesian Ones.

Subjects reported their beliefs about the probability of being in the high- detection-probability state after their enforcement experiences in part 2. Overall, 130 subjects were in state D, which means that they reported their beliefs after observing that the offender had been detected and sanctioned; 92 subjects were in state N and hence did not witness detection and sanctioning.[19]

The mean of the stated beliefs in state N is 38 percent, while the average stated belief in state D is 68 percent. In the Bayesian benchmark, these beliefs would be 36 percent in state N and 84 percent in state D. Thus, the average subject's estimate is very similar to that of the Bayesian benchmark in state N (p = .841 in a one-sample Wilcoxon signed-rank test against the benchmark).[20] The average subject's estimate in state D, however, clearly conveys the well-known phenomenon of conservatism, that is, an underweighting of the new information (p < .001 in a one- sample Wilcoxon signed-rank test; see, for example, Benjamin 2019).[21] Thus, we can(not) reject Bayesian updating for subjects in state D (N).

When we compare offenders and victims, we find that offenders' beliefs in state D and in state N are not statistically different from victims' beliefs in the respective states (Mann-Whitney *U*-tests [MWU]; offenders versus victims in state D: 70 percent versus 66 percent, p = .466; offenders versus victims in state N: 36 percent versus 39 percent, p = .357).

---

[18] Results are the same if we use a probit model.

[19] These numbers include subjects with firsthand enforcement experience and those who experienced enforcement by observing another pair. By design, the subjects in state N relied on firsthand enforcement experience, while some subjects in state D observed another pair.

[20] Throughout the paper, sign tests instead of Wilcoxon signed-rank tests yield qualitatively equivalent results for all comparisons of reported probabilities against the Bayesian benchmark.

[21] This pattern of conservatism in state D is robust if we consider only pairs with firsthand enforcement experience. The beliefs in state D differ significantly according to a Mann-Whitney *U*-test between pairs with taking and pairs without taking (72 percent [N = 76] versus 62 percent [N = 54], p = .004). Yet the 72 percent probability reported on average by the pairs with taking still represents conservatism.

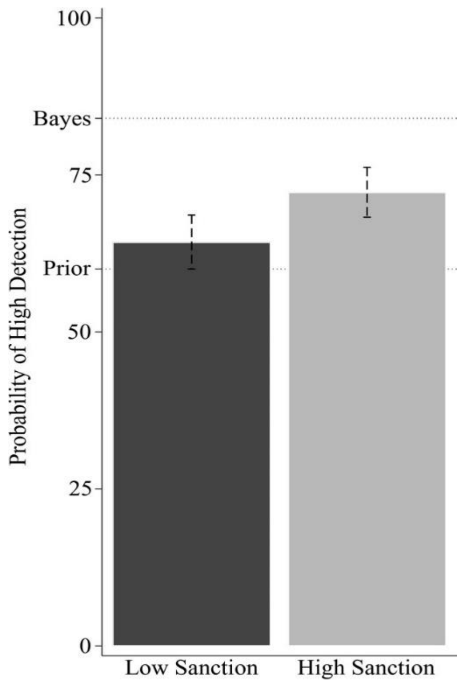

Figure 2. Conditional probability means in state D, with 95 percent confidence intervals (N = 130).

4.2.2. Sanctions' Severity and Learning about Enforcement Policy.

We start our inquiry with state D (see Figure 2). The mean belief of subjects with an enforcement experience involving a low sanction was 64 percent. This is significantly smaller than the mean belief of subjects with an enforcement experience involving a high sanction that amounted to 72 percent (MWU, p = .010). Thus, in our data, the severity of sanctions is an important moderator of how individuals update their beliefs about the probability of detection.

Result 1. In state D, the subjects who experienced a high sanction in part 1 predict a significantly higher probability of being in the high- detection-probability state than subjects who experienced a low sanction.

The experience of detection with a severe sanction does not induce an overshooting of the conditional probability. Beliefs for both sanction levels are significantly lower than the Bayesian benchmark of 84 percent (p < .001, one-sample Wilcoxon signed-rank test).

When we again explore a possible role dependency of beliefs, we find that offenders in state D report a belief amounting to 67 percent (73 percent) when they are subject to the low (high) sanction, whereas victims report 62 percent (71 percent). The differences in offenders' and victims' beliefs are not statistically significant for either sanction level (MWU, p = .48 and p = .811, respectively).

Turning to state N, we find that the mean belief by subjects with a low sanction in part 1 is 38 percent and therefore practically equal to the mean belief of subjects with a high sanction in part 1, amounting to 37 percent (MWU, p = .817). Hence, the severity of the sanction played no role in the updating of beliefs in state N.

Result 2. In state N, the subjects' beliefs about being in the high- detection-probability condition when the sanction in part 1 was high are comparable to those when it was low.

The significant influence of the sanction severity in state D is robust to the inclusion of the information elicited in part 4 as additional control variables in a regression analysis. Table 2 reports the results from ordinary least squares regressions and shows that the stated conditional probability increases by around 8 percentage points when the high sanction applies to the subject in question. As stated above, we find that the role in the pair is not influential.

Our analysis above relies on all observations, that is, subjects with and without direct experience of taking or enforcement. Our main result is robust to restricting the sample to subjects with firsthand enforcement experience in state D (although this reduces our sample to 76 observations). The average belief of subjects confronted with the high sanction is significantly higher than that of subjects confronted with the low sanction (MWU, S_L = 69 percent; S_H = 76 percent; p = .023). When we instead consider pairs without firsthand enforcement experience, the sample is reduced to 54 observations. At 67 percent, the belief that there is a high probability of detection is greater with a high sanction in part 1 than with a low sanction (58 percent), but the difference is not statistically significant (MWU, p = .15).

4.3. Part 3: Second Taking Decision

We find that the overall rate of taking in part 3 is similar to that in part 1 (77 and 76 percent, respectively). Remember that subjects were informed

Table 2. Conditional Probability of Subjects in State D

	(1)	(2)	(3)	(4)	(5)
High sanction	7.917**	8.250**	8.250**	7.566*	7.830*
	(2.960)	(2.842)	(2.838)	(3.088)	(3.187)
Taking		9.990**	9.990**	11.100**	9.318**
		(2.875)	(2.871)	(3.175)	(3.417)
Victim			−3.200	−2.817	−3.176
			(2.828)	(3.065)	(3.101)
Demographics	No	No	No	Yes	Yes
Questionnaires	No	No	No	No	Yes
Constant	64.3**	58.306**	59.906**	56.880**	60.857**
	(2.011)	(2.588)	(2.947)	(12.907)	(15.068)
R^2	.053	.135	.144	.186	.203

Note. Results are from ordinary least squares regressions. The dependent variable is the participants' reported probability that they are in the high detection state. High sanction is a dummy equal to one when the high sanction is applied. Taking is a dummy equal to one when the offender took money. Victim is a dummy equal to one for the role of the victim. Demographics include the participant's age, a gender dummy (male equals one), and field of study. Model 5 includes the cognitive reflection test score, the Big 5 measure for neuroticism (scale of 0 to 4), and the subject's response to the risk tolerance question (scale of 0 to 10) elicited in part 4. Standard errors are in parentheses. N = 130.

* $p < .05$.

** $p < .01$.

that the magnitude of the sanction would result randomly from the specified distribution, whereas the probability of detection in part 1 was carried over to part 3. As we established that the severity of the sanction was consequential for the beliefs about the relevant probability of detection, the severity in part 1 is expected to have an influence via the beliefs channel.

We report simple ordinary least squares regressions (Table 3) for taking in part 3.[22] The working sample consists of offenders in part 1 who experienced enforcement leading up to state D. We find that the stated belief in part 2 about the likelihood of being in the high-detection- probability state has a negative and significant relationship with the taking decision in part 3 (model 2). We do not find a

[22] These findings are robust to a change from the linear model to a probit model.

direct effect of sanction severity in part 1 on taking in part 3 (model 3). This finding is robust to the inclusion of the demographic variables (model 4), as the coefficients for the belief, taking behavior in part 1, and the attitude toward risk remain at least marginally significant.

Table 3. Taking in Part 3 by Offenders in State D in Part 1

	(1)	(2)	(3)	(4)
Belief Part 2		−.011**	−.010**	−.007+
		(.003)	(.003)	(.004)
Taking Part 1	.267*	.364**	.352**	.317*
	(.121)	(.116)	(.116)	(.127)
High Sanction			−.106	−.142
			(.102)	(.109)
Risk	.047+	.057*	.057*	.042+
	(.024)	(.022)	(.022)	(.024)
Demographics	No	No	No	Yes
Constant	.283*	.924**	.929**	.081
	(.118)	(.218)	(.218)	(.548)
R^2	.199	.327	.339	.448

Note. Results are from ordinary least squares regressions. The dependent variable is a dummy variable for taking in part 3. Belief Part 2 is the participants' reported probability in part 2 that they are in the high detection state. Taking Part 1 is a dummy equal to one when the subject took money in part 1. High Sanction is a dummy equal to one when the high sanction applied in part 1. Risk reflects the answer to the risk question elicited in part 4. Demographics include the subject's age, a gender dummy (male equals one), and field of study. Standard errors are in parentheses. $N = 65$.

+ $p < .10$.
* $p < .05$.
** $p < .01$.

5. Discussion

Our objective is to explore whether the severity of sanctions influences belief updating about the probability of detection. Our data analysis produces results 1 and 2, which show that the sanction level is in fact relevant in state D but not in state N.

While the irrelevance result for state N is in accordance with the Bayes's rule, the result for state D requires explanation. As a rationale for this finding, consider the following variation of the Bayes's rule formulation for the probability that the subject is in scenario H when part 1 ended in state D:

$$\pi(H|D, S_J) = \frac{\text{Prob}(H)[\alpha(D, S_J)\text{Prob}(D|H) + 1 - \alpha(D, S_J)]}{\{\text{Prob}(H)[\alpha(D, S_J)\text{Prob}(D|H) + 1 - \alpha(D, S_J)] + \text{Prob}(L)\,[\alpha(D, S_J)\text{Prob}(D|L) + 1 - \alpha(D, S_J)]\}}.$$

The weight $\alpha(D, S_J)$ is conditional on state D and the sanction severity S_J, where J = L, H and $S_L < S_H$. The weight $1 - \alpha(D, S_J)$ can be interpreted as a measure of conservatism bias. To see this, note that when $\alpha(D, S_J)$ is about 1, we obtain the adequate Bayesian posterior probability. In contrast, when $\alpha(D, S_J)$ is about 0, that is, when the enforcement experience that is informative about the true state is neglected, the belief is similar to the prior belief. Applying this framework to our setup, our finding that subjects with a low sanction in part 1 predict a probability of being in the high-detection-probability state lower than subjects with a high sanction results when $\alpha(D, S_L) < \alpha(D, S_H)$, that is, when subjects with the high sanction in part 1 take greater account of the detection signal. This is plausible for several reasons.

Because of the significant payoff implications, being detected when the high sanction applies represents a vivid enforcement experience and is likely to instill intensive emotions in subjects. Sanctioned offenders may be irritated and angry, while victims witnessing the sanctioning of their wrongdoer may feel happy (or at least satisfied). This arousal may induce subjects to incorporate the detection signal to a greater extent when they think about their conditional belief for being in the high-detection-probability state. In contrast, being detected when the sanction amounts only to 240 points (that is, only 10 percent of the points taken from the victim) is not a vivid experience for the average subject and will thus not induce a similar disproportionate weighting in subsequent judgments. The evidence in Kuhnen (2015) is consistent with this idea, as that study shows that individuals react more strongly to losses than to gains, which may be related to our distinction between negligible and sizable losses for the offenders in the detection state.

In the context of behavioral inattention, Gabaix (2019) refers to the distinction between system 1 and system 2 (popularized by Kahneman [2011]) to explain that system 1 unconsciously preselects which elements to bring to the attention of system 2 when making a decision. In line with our finding, having witnessed the enforcement experience involving the high sanction makes it more likely for system 1 to bring the detection experience from part 1 to the attention of system 2 when the decision about the belief about the high-detection-probability state is due.

There are also studies involving information costs that make the intuitive proposition of agents who tend to think more about aspects of the decision problem when more is at stake (for example, Caplin and Dean 2015). When subjects see that the detection state is associated with a high sanction, they are more likely to think that the probability of detection is an important piece of information and will thus search more actively and use information that is relevant to updating their beliefs. In summary, several lines of argument predict that the detection event will receive greater weight in the belief-updating process when it involves a greater sanction.

6. Conclusion

Using experimental data, we explore how individuals use enforcement experiences to update their beliefs about the probability of detection and whether the magnitude of the sanction plays a role in this process. We find that the stated belief after an offender's detection is significantly influenced by the magnitude of the sanction, in violation of the Bayesian prescription. In addition, we find that subjects who were not detected update their beliefs on average according to the prescription of Bayes's theorem, whereas subjects who were detected show conservative updating.

Our results are important from a policy-making perspective. We highlight that the widely held belief that the (perceived) probability of detection and the magnitude of the sanction are separate instruments cannot stand. Hence, a better understanding of the relationship between the two instruments is important for the law enforcement literature. Our results suggest that the severity of the sanction may have a double deterrent effect, since it may also affect people's perceived probability of detection. Thus, policy makers must bear in mind the underappreciated deterrence benefit of increasing sanctions when choosing the means to reach a target deterrence level and balancing other considerations against these benefits in setting regulatory schemes.

References

Akers, Ronald L., and Wesley G. Jennings. 2016. Social Learning Theory. Pp. 230–40 in *The Handbook of Criminological Theory*, edited by Alex R. Piquero. Chichester: John Wiley & Sons.

Andreoni, James. 1991. Reasonable Doubt and the Optimal Magnitude of Fines: Should the Penalty Fit the Crime? *Rand Journal of Economics* 22:385–95.

Anwar, Shamena, and Thomas A. Loughran. 2011. Testing a Bayesian Learning Theory of Deterrence among Serious Juvenile Offenders. *Criminology* 49:667– 98.

Apel, Robert. 2013. Sanctions, Perceptions, and Crime: Implications for Criminal Deterrence. *Journal of Quantitative Criminology* 29:67–101.

Barron, Kai. 2021. Belief Updating: Does the Good News Bad News Asymmetry Extend to Purely Financial Domains? *Experimental Economics* 24:31–58.

Baumann, Florian, and Tim Friehe. 2013. Cheap Talk about the Detection Probability. *International Game Theory Review* 15:52–67.

Bebchuk, Lucian Arye, and Louis Kaplow. 1992. Optimal Sanctions When Individuals Are Imperfectly Informed about the Probability of Apprehension. *Journal of Legal Studies* 21:365–70.

Becker, Gary S. 1968. Crime and Punishment: An Economic Approach. *Journal of Political Economy* 76:169–217.

Bénabou, Roland, and Jean Tirole. 2016. Mindful Economics: The Production, Consumption, and Value of Beliefs. *Journal of Economic Perspectives* 30(3):141–64.

Benjamin, Daniel. 2019. Errors in Probabilistic Reasoning and Judgment Bias. Pp. 2:69–186 in *Handbook of Behavioral Economics: Applications and Foundations,* edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson. Amsterdam: North-Holland.

Ben-Shahar, Omri. 1997. Playing without a Rulebook: Optimal Enforcement When Individuals Learn the Penalty Only by Committing the Crime. *International Review of Law and Economics* 17:409–21.

Bock, Olaf, Ingmar Baetge, and Andreas Nicklisch. 2014. hroot—Hamburg Registration and Organization Online Tool. *European Economic Review* 71:117– 20.

Brocas, Isabelle, and Juan D. Carrillo. 2012. From Perception to Action: An Economic Model of Brain Processes. *Games and Economic Behavior* 75:81–103.

Buser, Thomas, Leonie Gerhards, and Joel van der Weele. 2018. Responsiveness to Feedback as a Personal Trait. *Journal of Risk and Uncertainty* 56:165–92.

Caplin, Andrew, and Mark Dean. 2015. Revealed Preference, Rational Inattention, and Costly Information Acquisition. *American Economic Review* 105:2183–2203.

Charness, Gary, and Chetan Dave. 2017. Confirmation Bias with Motivated Beliefs. *Games and Economic Behavior* 104:1–23.

Cooter, Robert. 1998. Expressive Law and Economics. *Journal of Legal Studies* 27:585–607.

Costa, Paul T., and Robert R. McCrae. 1992. *Revised NEO Personality Inventory (NEO-PI-R) and NEO Five-Factor Inventory (NEO-FFI) Manual.* Lutz, FL: Psychological Assessment Resources.

Coutts, Alexander. 2019. Good News and Bad News Are Still News: Experimental Evidence on Belief Updating. *Experimental Economics* 22:369–95.

Daughety, Andrew, and Jennifer Reinganum. 2020. Reducing Unjust Convictions: Plea Bargaining, Trial, and Evidence Disclosure. *Journal of Law, Economics, and Organization* 36:378–414.

Deffains, Bruno, and Claude Fluet. 2020. Social Norms and Legal Design. *Journal of Law, Economics, and Organization* 36:139–69.

Deffin, Cyndie, and Maiva Ropaul. 2019. Subjective Probabilities of Sanction, Norms, Legitimacy and Everyday Life Crimes in Europe. *Revue d'économie politique* 129:143–68.

El-Gamal, Mahmoud A., and David M. Grether. 1995. Are People Bayesian? Uncovering Behavioral Strategies. *Journal of the American Statistical Association* 90:1137–45.

Engel, Christoph, and Daniel Nagin. 2015. Who Is Afraid of the Stick? Experimentally Testing the Deterrent Effect of Sanction Certainty. *Review of Behavioral Economics* 2:405–34.

Enke, Benjamin, Frederik Schwerter, and Florian Zimmermann. 2020. Associative Memory and Belief Formation. NBER Working Paper No. 26664. National Bureau of Economic Research, Cambridge, MA.

Ertac, Seda. 2011. Does Self-Relevance Affect Information Processing? Experimental Evidence on the Response to Performance and Non-performance Feedback. *Journal of Economic Behavior and Organization* 80:532–45.

Falk, Armin, Anke Becker, Thomas Dohmen, David Huffman, and Uwe Sunde. 2016. The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences. IZA Discussion Paper No. 9674. Institute of Labor Economics, Bonn.

———. 2018. Global Evidence on Economic Preferences. *Quarterly Journal of Economics* 133:1645–92.

Feess, Eberhard, Hannah Schildberg-Hörisch, Markus Schramm, and Ansgar Wohlschlegel. 2018. The Impact of Fine Size and Uncertainty on Punishment and Deterrence: Theory and Evidence from the Laboratory. *Journal of Economic Behavior and Organization* 149:58–73.

Fischbacher, Urs. 2007. z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics* 10:171–78.

Frederick, Shane. 2005. Cognitive Reflection and Decision Making. *Journal of Economic Perspectives* 19(4):25–42.

Gabaix, Xavier. 2019. Behavioral Inattention. Pp. 2:261–343 in *Handbook of Behavioral Economics: Applications and Foundations*, edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson. Amsterdam: North-Holland.

Garoupa, Nuno. 1997. The Theory of Optimal Law Enforcement. *Journal of Economic Surveys* 11:267–95.

Gennaioli, Nicola, and Andrei Shleifer. 2010. What Comes to Mind. *Quarterly Journal of Economics* 125:1399–1433.

Gotthard-Real, Alexander. 2017. Desirability and Information Processing: An Experimental Study. *Economics Letters* 152:96–99.

Grimm, Veronika, and Frederike Mengel. 2011. Let Me Sleep on It: Delay Reduces Rejection Rates in Ultimatum Games. *Economics Letters* 111:113–15.

Hay, Bruce L., and Kathryn E. Spier. 1997. Burdens of Proof in Civil Litigation: An Economic Perspective. *Journal of Legal Studies* 26:413–31.

Holt, Charles, and Angela M. Smith. 2009. An Update on Bayesian Updating. *Journal of Economic Behavior and Organization* 69:125–34.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus & Giroux.

Kaplow, Louis. 2011. On the Optimal Burden of Proof. *Journal of Political Economy* 119:1104–40.

Kelling, George E., and James Q. Wilson. 1982. The Police and Neighborhood Safety: Broken Windows. *Atlantic* 249(3):29–38.

Khadjavi, Menusch. 2015. On the Interaction of Deterrence and Emotions. *Journal of Law, Economics, and Organization* 31:287–319.

Köszegi, Botond. 2006. Ego Utility, Overconfidence, and Task Choice. *Journal of the European Economic Association* 4:673–707.

Kuhnen, Camelia M. 2015. Asymmetric Learning from Financial Information. *Journal of Finance* 70:2029–62.

Lieder, Falk, Thomas L. Griffiths, and Ming Hsu. 2018. Over-representation of Extreme Events in Decision-Making Reflects Rational Use of Cognitive Resources. *Psychological Review* 125:1–32.

Lochner, Lance. 2007. Individual Perceptions of the Criminal Justice System. *American Economic Review* 97:444–60.

Matsueda, Ross L., Derek A. Kreager, and David Huizinga. 2006. Deterring Delinquents: A Rational Choice Model of Theft and Violence. *American Sociological Review* 71:95–122.

Möbius, Markus M., Muriel Niederle, Paul Niehaus, and Tanya S. Rosenblat. 2022. Managing Self-Confidence: Theory and Experimental Evidence. *Management Science* 68:7793-7817.

Mullainathan, Sendhil. 2002. A Memory-Based Model of Bounded Rationality. *Quarterly Journal of Economics* 117:735–74.

Mungan, Murat C., and Jonathan Klick. 2016. Identifying Criminals' Risk Preferences. *Indiana Law Journal* 91:791–821.

Neo, Wei S., Michael Yu, Roberto A. Weber, and Cleotilde Gonzalez. 2013. The Effects of Time Delay in Reciprocity Games. *Journal of Economic Psychology* 34:20–35.

Polinsky, A. Mitchell, and Steven Shavell. 2007. The Theory of Public Enforcement of Law. Pp. 1:403–54 in *Handbook of Law and Economics*, edited by A. Mitchell Polinsky and Steven Shavell. Amsterdam: North-Holland.

Rammstedt, Beatrice, and Oliver P. John. 2007. Measuring Personality in One Minute or Less: A 10-Item Short Version of the Big Five Inventory in English and German. *Journal of Research in Personality* 41:203–12.

Rincke, Johannes, and Christian Traxler. 2011. Enforcement Spillovers. *Review of Economics and Statistics* 93:1224–34.

Rizzolli, Matteo, and Luca Stanca. 2012. Judicial Errors and Crime Deterrence: Theory and Experimental Evidence. *Journal of Law and Economics* 55:311– 38.

Sah, Raaj K. 1991. Social Osmosis and Patterns of Crime. *Journal of Political Economy* 99:1272–95.

Sanchirico, Chris W. 2012. Evidence: Theoretical Models. Pp. 203–70 in *Procedural Law and Economics*, vol. 8 of *Encyclopedia of Law and Economics*, edited by Chris William Sanchirico. Cheltenham: Edward Elgar.

Schildberg-Hörisch, Hannah, and Christina Strassmair. 2012. An Experimental Test of the Deterrence Hypothesis. *Journal of Law, Economics, and Organization* 28:447–59.

Schlag, Karl H., James Tremewan, and Joel van der Weele. 2015. A Penny for Your Thoughts: A Survey of Methods for Eliciting Beliefs. *Experimental Economics* 18:457–90.

Schotter, Andrew, and Isabel Trevino. 2014. Belief Elicitation in the Laboratory. *Annual Review of Economics* 6:103–28.

Sharot, Tali, and Neil Garrett. 2016. Forming Beliefs: Why Valence Matters. *Trends in Cognitive Sciences* 20:25–33.

Shavell, Steven. 2004. *Foundations of Economic Analysis of Law*. Cambridge, MA: Harvard University Press.

Toplak, Maggie E., Richard F. West, and Keith E. Stanovich. 2014. Assessing Miserly Information Processing: An Expansion of the Cognitive Reflection Test. *Thinking and Reasoning* 20:147–68.

Tversky, Amos, and Daniel Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science,* September 27, pp. 1124–31.

Zimmermann, Florian. 2020. The Dynamics of Motivated Beliefs. *American Economic Review* 110:337–36.