Selection Neglect in Policing Decisions

Juan B. González   Santiago De Martini   Ervyn Norza

Santiago M. Perez-Vincent *

June 5, 2025

(Most Recent Version Here)

Abstract

The wide racial disparities in policing decisions are often attributed to racial animus if they don't match accurate statistical predictions. But what if officers hold inaccurate beliefs? We study how selection neglect —not accounting for the data generating process when making predictions— can generate inaccurate statistical discrimination through self-reinforcing cycles of bias. In endogenous data environments, where predictions influence future data collection, neglecting selection causes decision makers to mislearn from feedback, perpetuating distorted beliefs. We design a novel framed field experiment to measure selection neglect and its consequences for discrimination among police officers in Latin America. Our design isolates how officers (mis)learn from selected data, how this biases policing decisions, and how it distorts the value placed on unbiased information. By estimating individual-level biases and linking it to discriminatory behavior, we show how statistical discrimination can emerge from cognitive error rather than animus, and how it may persist even under data-driven policing.

*González: Department of Economics, University of Southern California, juanbgon@usc.edu. De Martini: Department of Economics, University of Southern California, sd59576@usc.edu. Norza: National Police of Colombia, eh.norza10@uniandes.edu.co. Perez-Vincent: Inter American Development Bank, santiagoper@iadb.org. We thank Isabelle Brocas, Colin Camerer, Juan Carrillo, Giorgio Coricelli, Jonathan Libgober, Alejandro Martínez-Marquina, Paulina Oliva, Emanuel Vespa, and the participants of the Graduate Behavioral and Experimental Economics Conference at Caltech for thoughtful feedback on the study.

There are systematic racial disparities in the criminal justice system. Police officers and judges are more likely to stop, search, fine, use force, charge, detain, and incarcerate civilians who belong to a racial minority group or have lower socioeconomic status [1]. While these disparities are well-documented, their origin remains contested: are they driven by animus or by statistical predictions? Much of the literature assumes that if disparities cannot be justified by accurate statistical discrimination, they must stem from taste-based bias. But this reasoning overlooks a crucial possibility: officers and judges may hold inaccurate beliefs. Distinguishing between discrimination rooted in biased preferences and that arising from biased inference is essential for designing effective interventions. Whereas animus is difficult to change, cognitive distortions and inaccurate beliefs can often be corrected through targeted information or training (Bohren et al., 2019; Dube et al., 2024; Paluck & Green, 2009).

A potential cause of inaccurate statistical discrimination is selection neglect —a cognitive bias in which individuals don't account for data selection when making predictions from this data (Enke, 2020).[2] In endogenous data environments, where predictions influence data collection, selection neglect creates self-reinforcing cycles of bias: biased predictions lead to skewed sampling, which in turn reinforces prediction errors (Esponda & Vespa, 2018). Policing decisions offer a textbook case of an endogenous data environment. Officers rely on past crime data —such as arrests— to guide crime predictions and policing decisions[3], with the goal of maximizing arrests (Feigenberg & Miller, 2025; Stashko, 2023). But crime data is itself shaped by previous policing decisions: conditional on underlying crime levels, more arrests occur in more heavily policed areas (Chen et al., 2023). If officers neglect this selection process, they may incorrectly infer that over-policed areas are more criminal, reinforcing patrol allocations and perpetuating overpolicing. Even in the absence of animus, selection neglect can create

[1] For recent evidence, see Abrams et al. (2012), Aggarwal et al. (2025), Arnold et al. (2018), Ba et al. (2021), Chen et al. (2023), Edwards et al. (2019), Feigenberg and Miller (2025), Fryer (2019), Goncalves and Mello (2021), Gupta et al. (2016), Hoekstra and Sloan (2022), Knox et al. (2020), Pierson et al. (2020), Rizzotti (2024), and Rozema and Schanzenbach (2019).

[2] As a widespread bias among the population, selection neglect has been used to explain echo chambers (Brundage et al., 2024), overoptimism in investment (Barron et al., 2024; Jehiel, 2018), and distorted views about income distributions (Cruces et al., 2013), for example.

[3] Predictive policing algorithms are increasingly used alongside officer judgment (Mohler et al., 2015), but they too are vulnerable to selection neglect and thus amplify racial bias (Brayne, 2020; Lum & Isaac, 2016).

cycles of bias that mirror the empirical patterns of overpolicing. While the role of selection neglect in biased crime predictions been formalized theoretically (Hübert & Little, 2023), empirical evidence remains scarce due to the inherent challenges of disentangling cause and effect in inherently endogenous and opaque institutional settings.

In this paper, we design a framed field experiment to measure selection neglect and test whether it distorts belief updating, information acquisition, and policing decisions in ways that generate statistical discrimination. Participants complete the *Patrolling Task*, an abstract setting involving three hypothetical "neighborhoods" with varying levels of "criminality." Neighborhoods may or may not be patrolled; if patrolled, a crime is detected with a probability equal to the area's criminality level, and if not, no crime is observed. To isolate inference from taste-based bias, participants receive no identifying information about neighborhoods, only patrol and crime outcomes. The task is implemented in three sequential blocks. In Block 1, all neighborhoods are patrolled across 15 rounds, generating an unselected sample of crime outcomes. Participants observe these signals and are incentivized to accurately rank neighborhoods by criminality, allowing us to estimate individual updating styles relative to the Bayesian benchmark. In Block 2, patrols are randomly assigned to one or two neighborhoods per round, introducing exogenous selection. Holding individual updating fixed, we measure selection neglect by comparing belief updates after informative (patrolled) versus uninformative (unpatrolled) null signals. Block 3 introduces endogenous selection: over 40 trials, participants now choose which neighborhood to patrol and are rewarded for the crimes caught. This multi-armed bandit design lets us study how selection neglect interacts with information acquisition. Finally, participants report their willingness to pay for learning the true level of criminality, separately for each neighborhood, before choosing one to patrol for a final set of rounds. This elicitation captures how selection neglect affects the perceived value of accurate information.

This paper contributes to a large applied literature documenting discrimination in the criminal justice system. Within this literature, several recent studies suggest that these disparities can arise from statistical predictions. For example, Feigenberg and Miller (2025) suggest that class disparities in traffic stops and searches are driven by officers' higher expected litigation costs when stopping high-income drivers. Similarly, Holz et al. (2023) show that officers become more likely to use force after a colleague is injured, consistent with updates in perceived risk. However, a common assumption in the literature is that if observed disparities deviate from *accurate* statistical predictions, they must be driven by taste-based discrimination. This overlooks the possibility that agents may make systematically *inaccurate* predictions due to cognitive distortions. Indeed, recent evidence shows that even high-stakes decision-makers, such as judges, deviate from Bayesian updating by overreacting to signals (Bhuller & Sigstad, 2024). Recognizing the role of learning distortions is crucial, as cognitive biases are more amenable to intervention than preferences or animus. For instance, promoting reflective thinking (Dube et al., 2024) or procedural reasoning (Owens et al., 2018) has been shown to reduce discrimination in policing. We contribute to this agenda by experimentally measuring how selection neglect distorts learning and decisions in a setting that mimics the key features of policing decisions.

A growing experimental literature studies how individuals learn under misspecified models, often leading to systematic errors in belief formation (Bohren & Hauser, 2025). These errors frequently arise from distorted attention to or misinterpretation of information (Bordalo et al., 2023; Fréchette et al., 2024), and can generate learning traps and persistent misbeliefs (Esponda et al., 2024; Gagnon-Bartsch et al., 2023). We contribute to this literature by showing how a specific cognitive bias can distort belief updating in environments where the data-generating process is itself shaped by prior predictions. More broadly, we relate to research on how cognitive heuristics, such as stereotypical reasoning, amplify bias in inference and decision-making (Bordalo et al., 2016; Esponda et al., 2023). Finally, we connect to work on the valuation of information, showing how belief distortions can affect not only predictions, but also how decision-makers value access to unbiased signals (Afrouzi et al., 2023; Ambuehl & Li, 2018; Charness et al., 2021).

Finally, our work builds on recent advances in understanding how individuals trade off exploration and exploitation when acquiring and acting on reward-relevant information. These dynamics and the biases that shape them have been extensively studied in cognitive science (see Palminteri and Lebreton (2022) for a review), with recent work focusing on how to disentangle exploration from exploitation behavior (Lizzeri et al., 2024). These insights have been applied to practical settings such as research funding allocation (Zhuo, 2023) and product selection in digital platforms (Jin et al., 2021). In the context of policing, Che et al. (2024) model "greedy" officers who dismiss the future value of exploration, leading to over-exploitation and persistent over-policing. We extend this research by providing experimental evidence on how selection neglect distorts exploration incentives and decisions. Our design allows us to separately identify exploitation, exploration, and the perceived value of unbiased information in a setting that mirrors key features of real-world policing.

We structure the paper as follows. Section 1 describes the experimental design and its implementation. Section 2 formalizes how selection neglect distorts belief updating and generates discrimination in endogenous data environments, providing a framework for our empirical analysis.

1 Experimental Design

The Patrolling Task

The Patrolling Task models an endogenous data environment where participants make predictions and exploration decisions upon observing potentially selected signals. We present participants with three boxes that represent *neighborhoods* that can be patrolled to catch the crimes happening there. Each neighborhood n has a level of *criminality* (c_n) — a constant probability of catching a crime there if the neighborhood is patrolled. Importantly, participants don't know these parameters. In contrast, if a neighborhood is not patrolled, no crimes are caught there. We train participants to understand this setting and its data generating process. Additionally, we prevent memory issues by providing access to the entire history of patrols and crime realizations at any time, using a mouse-tracking design[4] that records whether participants request information about patrols, crimes caught, or both. Overall, this setting subtly reframes the classic multi-arm bandit problem[5], which we separate into a

[4] We ask participants to click in separate buttons to reveal each part of the history, and record the buttons and the order of clicks. This method has been previously successful in providing direct evidence of what information is considered when making decisions (De Martini et al., 2025).

[5] In the standard multi-arm bandit setting, a decision maker faces several slot machines (*one-arm bandits*) with different probability of reward. The agent pulls one arm in each trial and observes whether this resulted in a reward. At each pull decision, the agent must trade off the incentive to keep pulling the arm that has

3-block design to measure selection neglect and disentangle exploration from exploitation.

Figure 1: Patrolling Task – Block 2

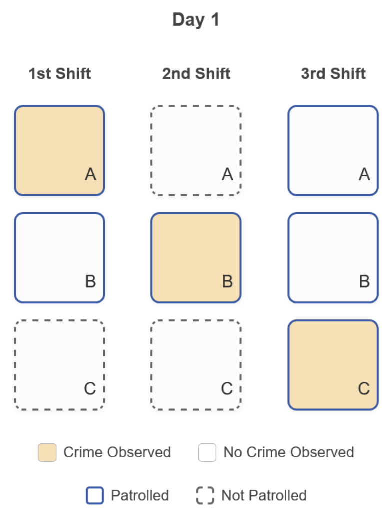

In Block 1, participants are introduced to 3 neighborhoods (A, B, C), each neighborhood with a different criminality level [6], without replacement. The goal of this block is to accurately infer these criminality levels from crime signals —whether crime was caught or not, which are the outcome of patrolling. During 9 rounds, participant observe whether each neighborhood was patrolled and the resulting crime signal. Importantly, during Block 1 every neighborhood is always patrolled, so crime signals are *not selected*. To provide more learning opportunities, we provide 3 signals of each neighborhood at each round, which we frame as 3 patrolling shifts. Figure XXX displays the screen for Block 1, with the 9 signals participant observe at each round. After observing these signals, participants report their belief of criminality for each neighborhood, using 3 sliders from 0 to 100%, and are paid according to their accuracy. These predictions don't have any effect on patrolling decisions, as all neighborhoods are always patrolled. We use the 27 belief elicitations from Block 1 to pin down how each participant updates their beliefs upon observing unselected signals, providing a benchmark for next blocks.

Block 2 introduces *exogenous selection* by having neighborhoods randomly patrolled, so

provided rewards more often —*exploiting* the current information, and the incentive to pull other arms that could potentially be more profitable — *exploring* to gather new information.

[6]Criminality levels are randomly selected among the set $\{1/9, 2/9, 3/9, 4/9, 5/9, 6/9, 7/9, 8/9\}$

patrolling is orthogonal to criminality predictions. Participants are now introduced to 3 neighborhoods (D,E,F), which are different from the ones in Block 1, so learning cannot be carried across blocks. As in the previous block, at each round participants observe which neighborhoods are patrolled during 3 shifts and the resulting 3 crime signals from each neighborhood. They can use these 9 signals to update and report their believed criminality levels, and are incentivized for accuracy, during 12 rounds. Unlike in Block 1, neighborhoods are not always patrolled, so crime data is selected and there is room for selection neglect to distort predictions. Participants who neglect selection will update in a similar way following an informative signal of no crime when a neighborhood is patrolled and following an uninformative signal of no crime when a neighborhood is not patrolled. We use the 36 predictions of criminality of each participant to measure selection neglect at the individual level.

To better understand the role of differential selection, we constraint the randomization of criminality across neighborhoods in Block 2. We partition the parameter space of crime probabilities into $\{Low, Mid, High\}$, and randomly assign each neighborhood to one of these partitions [7]. We randomize participants into 3 treatments. As a benchmark, in the *Orthogonal* treatment each neighborhood is patrolled 18 times. In contrast, in the *Predictive Policing* treatment, the $\{Low, Mid, High\}$ neighborhoods are patrolled $\{9, 18, 27\}$ times, respectively. This mirrors the real setting, where high crime areas are more intensely policed, but keeping patrolling decisions separate from crime predictions. Finally, the *Anti-predictive* treatment reverses the relation between crime and patrolling, having the $\{Low, Mid, High\}$ neighborhoods patrolled $\{27, 18, 9\}$ times, respectively. We use the variation in selection from these treatments to test whether selection neglect generates systematically biased predictions under biased data selection.

[7] In particular, the partition is $\{\{1/9, 2/9, 3/9\}, \{4/9, 5/9\}, \{6/9, 7/9, 8/9\}\}$. Thus, the neighborhood that is assigned a *Low* criminality will have a crime probability randomly chosen among $\{1/9, 2/9, 3/9\}$.

Figure 2: Patrolling Task – Block 3

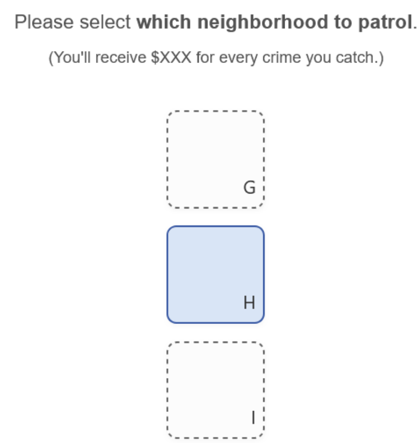

Block 3 closely follows a multi-arm bandit design, where participants are now incentivized to observe crimes rather than to predict criminality. They start having no information about 3 new neighborhoods (G,H,I), and choose at each round which of them to patrol. As a result of their patrolling decision, they observe whether crime was observed there or not, and are paid according to the number of crimes they observe. In each of 45 rounds, they make a new patrolling decision and observe the resulting crime signals. At each decision participants face the trade-off between patrolling the neighborhood they perceive to be more criminal (*exploiting*), and gathering new information (*exploring*). As we can estimate the beliefs participants have at each trial from how they updated in the first two blocks, we can clearly disentangle exploitation from exploration decisions. This design allows to test whether selection neglect biases exploration decisions, creating biases not only in predictions but in policing decisions.

Value of unbiased information

We elicit the willingness to pay for unbiased information by the end of the 45 trials of Block 3. Participants learn they can choose only one neighborhood to patrol for the next 45 rounds. To guide their decision, they can get a precise estimate of criminality of one neighborhood. We use the BDM mechanism to elicit their maximum willingness to pay for this estimate independently for each neighborhood. After the elicitation, we randomly select a price and a neighborhood, determining which information participants observe. Then, participants choose which neighborhood to patrol and crime outcomes are realized. This elicitation serves multiple purposes. First, the overall valuation of unbiased information is additional evidence of the value subjects place in exploration. Second, if selection neglect affects the value of exploration, we expect subjects who neglect selection to undervalue information about unexplored neighborhoods. Third, as we can estimate subjects expectations from the first two blocks, we can decompose deviations from the optimal Bayesian valuation into three distortions: generated by non-Bayesian update, generated by selection neglect, and a third component that is not explained by predictions.

Cognitive Skills

Finally, we test whether numeracy skills and cognitive reflection are associated with selection neglect. We include two questions to test numeracy skills (Kahan et al., 2012), and a variation of the standard cognitive reflection questions (Frederick, 2005). The questions are available in the Appendix A.

2 Conceptual Framework and Analysis

We develop a simple formal framework to fix ideas about the experimental design and guide a model-based empirical analysis [8].

Selection Neglect in Endogenous Data Environments

A decision maker is presented with N neighborhoods. Each neighborhood n has some constant criminality level c_n, which represents the constant probability of a crime happening in the area. At each time period t, whether a crime c_{nt} happens in neighborhood n is the outcome of a Bernoulli process with parameter c_n. Importantly, the outcomes are independent across neighborhoods and time. To observe and catch a crime c_{nt}, the neighborhood has to be patrolled: $p_{nt} = 1$. We define an arrest a_{nt} as the event where a crime happens and is caught, so $a_{nt} = c_{nt}p_{nt}$. We assume the decision maker has a complete information set that comprises

[8] For more general theoretical approaches to selection neglect and predictive policing, see Che et al. (2024) and Hübert and Little (2023)

all past patrolling and arrest outcomes. For a given neighborhood, this information set can be characterized through the total number of patrols $P_{nt} = \sum^t p_{nt}$, and the frequency with which those patrols resulted in arrests: $f_{nt} = \frac{\sum^t a_{nt}}{P_{nt}}$.

We assume the agent is uncertain about the criminality c_n. Their prior can be characterized by a Beta distribution governed by parameters c_n^0 and ω_n, such that

$$E[c_n|c_n^0, \omega_n] = c_n^0 \quad \text{and} \quad \text{var}(c_n|c_n^0, \omega_n) = \frac{c_n^0(1-c_n^0)}{\omega_n + 1}.$$

In this formulation of the Beta distribution[9], c_n^0 represents the prior belief (i.e., prior mean) of criminality c_n, and ω_n the weight or confidence put in this prior, which will affect future updating. Thus, this characterization models deviations from Bayesian updating through inaccurate priors and under or over reactions to signals. To introduce selection neglect in this setting, note that an agent that completely neglects data selection interprets the lack of signals (no patrolling) as a signal of no crime, as it doesn't account for the dependency of crime data on patrolling. To keep it simple, we model selection neglect by assuming the agent confounds the total number of patrols with the total number of periods, $\tilde{P}_{nt} = P_{nt} + \lambda(t - P_{nt})$, where λ measures the extent of the neglect. Thus, the posterior belief $\tilde{F}_{nt} \equiv \tilde{F}_t(c_n|\tilde{P}_{nt}, \tilde{f}_{nt})$ is characterized by a Beta distribution with modified parameters $\tilde{c}_{nt}$ and $\tilde{\omega}_{nt}$, such that

$$\tilde{c}_{nt} = \frac{\omega_n}{\tilde{\omega}_{nt}} c_n^0 + \left(1 - \frac{\omega_n}{\tilde{\omega}_{nt}}\right) \tilde{f}_{nt} \quad \text{and} \quad \tilde{\omega}_{nt} = \omega_n + \tilde{P}_{nt}. \tag{1}$$

Selection neglect distorts predictions of criminality by concentrating and shifting towards zero the posterior beliefs of criminality of less patrolled neighborhoods. Formally, $\tilde{c}_{nt}$ is weakly decreasing in the extent of selection neglect λ, as well as the variance of the posterior belief. Intuitively, because a neglecting agent overlooks the dependency of crime signals on patrolling, they confidently believe an under-policed neighborhood to be less criminal. Thus, selection neglect can generate *inaccurate* statistical discrimination by mistakenly reducing the believed criminality of under-policed areas relative to over-policed ones.

The direct bias of selection neglect on predictions compounds in endogenous data environments where predictions guide further data collection — patrolling decisions. We model

[9] The standard formulation of the Beta distribution is governed by parameter α, β, such that $E[c|\alpha, \beta] = \frac{\alpha}{\alpha+\beta}$. We modify this formulation by defining $c^0 = \frac{\alpha}{\alpha+\beta}$ and $\omega = \alpha + \beta$ to ease interpretation.

patrolling decisions by assuming the decision maker maximizes an utility function $\pi_{nt}(a_{nt}, \theta_n)$ that is increasing in the number of arrests, and that can depend on some taste parameter θ_n. When deciding where to patrol, the agent trades off two different incentives. On one hand, they have an incentive to *exploit* the information they have currently gathered and patrol the neighborhood with the highest expected utility. On the other hand, there is some future value in *exploring* and collecting information about less-known neighborhoods. By combining both incentives into one value index, we can characterize the decision of where to patrol as choosing the neighborhood with the highest value index V_{nt}, which we assume to have an UCB[10] form:

$$V_{nt} = \int \pi_{nt} \; d\tilde{F}_{nt} + \sqrt{\frac{\gamma}{\tilde{P}_{nt}}} \tag{2}$$

The first component of V_{nt} captures the expected value of patrolling, this is, the value of exploiting the information already available. Then, an exploration bonus captures the value of gathering new information, which increases in the relative taste for exploration γ and decreases in the *perceived* number of explorations already conducted. Selection neglect can amplify this number for less patrolled neighborhoods, since $\tilde{P}_{nt}$ is weakly increasing in the extent of neglect λ, reducing their perceived value of exploration.

In conclusion, selection neglect can thus bias predictions and decisions in endogenous data environments through the two channels represented in Equation (2). First, selection neglect distorts predictions of criminality and reduces the relative expected value of patrolling under-policed neighborhoods. Second, it conflates exploration values across neighborhoods, which should be higher for less patrolled areas if selection was accounted for. Through these two channels, selection neglect can generate cycles of bias where under-policed neighborhoods become increasingly less attractive to patrol, pushing officers to redirect resources to traditionally policed areas instead.

[10]Although many index solutions have been proposed to dynamic programming problems in multi-arm bandits (see Gittins (1979) and Lattimore and Szepesvári (2020)), the Uppper-Confidence Bounds (UCB) approximation is robust to many bandit specifications and highly convenient for estimation, making it the most common in applied research (Jin et al., 2021; Zhuo, 2023). Nevertheless, our argument for the effects of selection neglect doesn't depend on the functional form of the index approximation as long as the exploration bonus is decreasing in the total number of previous explorations.

Value of Information

After Block 3, participants report their maximum willingness to pay for unbiased information about the criminality of each neighborhood, this is, c_n. To characterize a benchmark for this value, we first need to note that this information is only valuable insofar it drives the agent to change their decision of which neighborhood to patrol for the last 20 trials. As there are is no exploration value from this last patrolling decision —the information you'll gather will not be useful when the block ends, the agent should pick the neighborhood with the highest expected criminality. Taking this into account, let $k = \text{argmax}_{j \neq n} \tilde{c}_{jt}$ be the neighborhood with the highest expected criminality apart from n. Then, the ex-ante expected value of the patrolling decision conditional on the unbiased information c_n is given by:

$$\tilde{F}_{nt}(\tilde{c}_{kt})\tilde{c}_{kt} + (1 - \tilde{F}_{nt}(\tilde{c}_{kt})) \int_{\tilde{c}_k}^{1} c_n \, d\tilde{F}_{nt} \tag{3}$$

The first component of Equation 3 represents the expected value of the decision if the information about n makes it a suboptimal choice, whereas the second component reflects it when c_n ends up being higher than the other expected criminalities so n is the optimal choice. To isolate the value of information I_{nt}, we subtract the expected value of the choice without information (choosing the highest $\tilde{c}_{nt}$) to Equation 3.

Empirical Analysis

The block design of the patrolling task closely follows the steps on which the conceptual model is built. In Block 1, there is no selection ($\tilde{P}_{nt} = t \; \forall \; n, \lambda$), so we can analyze signals and elicited beliefs to understand individuals' updating style. We fix participants priors to $c_n^0 = 1/2 \; \forall \; n$. As prior beliefs should be then identical across neighborhoods, we assume a participant to be equally confident about them and we focus on estimating $\omega_n = \omega \; \forall \; n$, a single parameter governing each individual's updating. Having pinned down how much each participant weights prior and signals, we carry the estimated parameter ω into Block 2. Block 2 introduces differential selection, leaving room for selection neglect to distort belief updating. We use the trials from this block to estimate λ, the extent to which individual neglects selection.

Next, we use the fitted parameters (ω, λ) to estimate the different components of Equation

(2). We use the information from the first two blocks to characterize the expected value of exploitation for each neighborhood at each period. We are left with estimating the individual preference for exploration γ, which would explain non expected value maximizing decisions.

Finally, to analyze the elicited value of information, we characterize three benchmarks for the information value using (3). First, we calculate the posterior beliefs $\tilde{F}_{nt}$ using all the information from Blocks 1 and 2, this is, considering both non-Bayesian updating and selection neglect. We denote the information value according to this posterior by I^{Full}. Second, we calculate the posterior assuming there is no selection neglect but maintaining the individually estimated updating, which results in information value I^{NB}. Finally, we calculate the information value corresponding to the correct Bayesian posteriors, I^{Bay}. Altogether, we use these three benchmarks to measure how selection neglect biases information value, and to test which benchmark is closer to the elicited willingness to pay.

References

Abrams, D. S., Bertrand, M., & Mullainathan, S. (2012). Do Judges Vary in Their Treatment of Race? *The Journal of Legal Studies*, *41*(2), 347–383.

Afrouzi, H., Kwon, S. Y., Landier, A., Ma, Y., & Thesmar, D. (2023). Overreaction in Expectations: Evidence and Theory. *The Quarterly Journal of Economics*, *138*(3), 1713–1764.

Aggarwal, P., Brandon, A., Goldszmidt, A., Holz, J., List, J. A., Muir, I., Sun, G., & Yu, T. (2025). High-frequency location data show that race affects citations and fines for speeding. *Science*, *387*(6741), 1397–1401.

Ambuehl, S., & Li, S. (2018). Belief updating and the demand for information. *Games and Economic Behavior*, *109*, 21–39.

Arnold, D., Dobbie, W., & Yang, C. S. (2018). Racial Bias in Bail Decisions. *The Quarterly Journal of Economics*, *133*(4), 1885–1932.

Ba, B., Knox, D., Mummolo, J., & Rivera, R. (2021). The role of officer race and gender in police-civilian interactions in Chicago. *Science*, *371*(6530), 696–702.

Barron, K., Huck, S., & Jehiel, P. (2024). Everyday Econometricians: Selection Neglect and Overoptimism When Learning from Others. *American Economic Journal: Microeconomics*, *16*(3), 162–198.

Bhuller, M., & Sigstad, H. (2024). Feedback and Learning: The Causal Effects of Reversals on Judicial Decision-Making. *The Review of Economic Studies*, rdae073.

Bohren, J. A., & Hauser, D. N. (2025). Misspecified Models in Learning and Games.

Bohren, J. A., Imas, A., & Rosenberg, M. (2019). The Dynamics of Discrimination: Theory and Evidence. *American Economic Review*, *109*(10), 3395–3436.

Bordalo, P., Coffman, K., Gennaioli, N., & Shleifer, A. (2016). Stereotypes. *The Quarterly Journal of Economics*, *131*(4), 1753–1794.

Bordalo, P., Conlon, J. J., Gennaioli, N., Kwon, S. Y., & Shleifer, A. (2023). How People Use Statistics. *NBER Working Paper*.

Brayne, S. (2020, November). *Predict and Surveil: Data, Discretion, and the Future of Policing*. Oxford University Press.

Brundage, M., Little, A. T., & You, S. (2024). Selection Neglect and Political Beliefs. *Annual Review of Political Science*, *27*(Volume 27, 2024), 63–85.

Charness, G., Oprea, R., & Yuksel, S. (2021). How do People Choose Between Biased Information Sources? Evidence from a Laboratory Experiment. *Journal of the European Economic Association*, *19*(3), 1656–1691.

Che, Y.-K., Kim, J., & Mierendorff, K. (2024, September). Predictive Enforcement.

Chen, M. K., Christensen, K. L., John, E., Owens, E., & Zhuo, Y. (2023). Smartphone Data Reveal Neighborhood-Level Racial Disparities in Police Presence. *The Review of Economics and Statistics*, 1–29.

Cruces, G., Perez-Truglia, R., & Tetaz, M. (2013). Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *Journal of Public Economics*, *98*, 100–112.

De Martini, S., González, J. B., & Perez-Vincent, S. M. (2025). The Impact of Crime on Public Transport Demand: Evidence from Six Latin American Capitals.

Dube, O., MacArthur, S. J., & Shah, A. K. (2024). A Cognitive View of Policing. *The Quarterly Journal of Economics*.

Edwards, F., Lee, H., & Esposito, M. (2019). Risk of being killed by police use of force in the United States by age, race–ethnicity, and sex. *Proceedings of the National Academy of Sciences*, *116*(34), 16793–16798.

Enke, B. (2020). What You See Is All There Is. *The Quarterly Journal of Economics*, *135*(3), 1363–1398.

Esponda, I., Oprea, R., & Yuksel, S. (2023). Seeing What is Representative. *The Quarterly Journal of Economics*, *138*(4), 2607–2657.

Esponda, I., & Vespa, E. (2018). Endogenous sample selection: A laboratory study. *Quantitative Economics*, *9*(1), 183–216.

Esponda, I., Vespa, E., & Yuksel, S. (2024). Mental Models and Learning: The Case of Base-Rate Neglect. *American Economic Review*, *114*(3), 752–782.

Feigenberg, B., & Miller, C. (2025, March). Class Disparities and Discrimination in Traffic Stops and Searches.

Fréchette, G. R., Vespa, E., & Yuksel, S. (2024, November). Extracting Models From Data Sets: An Experiment.

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

Fryer, R. G. (2019). An Empirical Analysis of Racial Differences in Police Use of Force. *Journal of Political Economy*, *127*(3), 1210–1261.

Gagnon-Bartsch, T., Rabin, M., & Schwartzstein, J. (2023). Channeled Attention and Stable Errors.

Gittins, J. C. (1979). Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society: Series B (Methodological)*, *41*(2), 148–164.

Goncalves, F., & Mello, S. (2021). A Few Bad Apples? Racial Bias in Policing. *American Economic Review*, *111*(5), 1406–1441.

Gupta, A., Hansman, C., & Frenchman, E. (2016). The Heavy Costs of High Bail: Evidence from Judge Randomization. *The Journal of Legal Studies*, *45*(2), 471–505.

Hoekstra, M., & Sloan, C. (2022). Does Race Matter for Police Use of Force? Evidence from 911 Calls. *American Economic Review*, *112*(3), 827–860.

Holz, J. E., Rivera, R. G., & Ba, B. (2023). Peer Effects in Police Use of Force. *American Economic Journal: Economic Policy*, *15*(2), 256–291.

Hübert, R., & Little, A. T. (2023). A Behavioural Theory of Discrimination in Policing. *The Economic Journal*, *133*(655), 2828–2843.

Jehiel, P. (2018). Investment Strategy and Selection Bias: An Equilibrium Perspective on Overoptimism. *American Economic Review*, *108*(6), 1582–1597.

Jin, R., Simchi-Levi, D., Wang, L., Wang, X., & Yang, S. (2021). Shrinking the Upper Confidence Bound: A Dynamic Product Selection Problem for Urban Warehouses. *Management Science*, *67*(8), 4756–4771.

Kahan, D. M., Peters, E., Wittlin, M., Slovic, P., Ouellette, L. L., Braman, D., & Mandel, G. (2012). The polarizing impact of science literacy and numeracy on perceived climate change risks. *Nature Climate Change*, *2*(10), 732–735.

Knox, D., Lowe, W., & Mummolo, J. (2020). Administrative Records Mask Racially Biased Policing. *American Political Science Review*, *114*(3), 619–637.

Lattimore, T., & Szepesvári, C. (2020, September). *Bandit Algorithms* (1st edition). Cambridge University Press.

Lizzeri, A., Shmaya, E., & Yariv, L. (2024, May). Disentangling Exploration from Exploitation.

Lum, K., & Isaac, W. (2016). To Predict and Serve? *Significance*, *13*(5), 14–19.

Mohler, G. O., Short, M. B., Malinowski, S., Johnson, M., Tita, G. E., Bertozzi, A. L., & Brantingham, P. J. (2015). Randomized Controlled Field Trials of Predictive Policing. *Journal of the American Statistical Association*, *110*(512), 1399–1411.

Owens, E., Weisburd, D., Amendola, K. L., & Alpert, G. P. (2018). Can You Build a Better Cop? *Criminology & Public Policy*, *17*(1), 41–87.

Palminteri, S., & Lebreton, M. (2022). The computational roots of positivity and confirmation biases in reinforcement learning. *Trends in Cognitive Sciences*, *26*(7), 607–621.

Paluck, E. L., & Green, D. P. (2009). Prejudice reduction: What works? A review and assessment of research and practice. *Annual Review of Psychology*, *60*, 339–367.

Pierson, E., Simoiu, C., Overgoor, J., Corbett-Davies, S., Jenson, D., Shoemaker, A., Ramachandran, V., Barghouty, P., Phillips, C., Shroff, R., & Goel, S. (2020). A large-scale analysis of racial disparities in police stops across the United States. *Nature Human Behaviour*, *4*(7), 736–745.

Rizzotti, L. (2024). Racial Discrimination against Crime Victims: Evidence from Violent Crime Investigations.

Rozema, K., & Schanzenbach, M. (2019). Good Cop, Bad Cop: Using Civilian Allegations to Predict Police Misconduct. *American Economic Journal: Economic Policy*, *11*(2), 225–268.

Stashko, A. (2023). Do Police Maximize Arrests or Minimize Crime? Evidence from Racial Profiling in U.S. Cities. *Journal of the European Economic Association*, *21*(1), 167–214.

Zhuo, R. (2023). Exploit or explore? An empirical study of resource allocation in research labs.

A Numeracy Questions

1. **Numeracy 1**: The chance of getting a viral infection is .0005. Out of 10,000 people, about how many of them are expected to get infected?

2. **Numeracy 2**:Imagine that we roll a fair, six-sided die 1,000 times. Out of 1,000 rolls, how many times do you think the die would come up as an even number?

3. **Cross-multiplication 1**: In a jar there are some balls. If 75% of the jar has 60 balls, how many balls has the jar?

4. **Cross-multiplication 2**: 2 out of 3 students in a class are right handed. If there are 18 right handed students in the class, how many students are there in the class?

5. **CRT**: In a lake, there is a patch of lilypads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake?