# INATTENTIVE INFERENCE

**Thomas Graeber**
Harvard Business School, USA

**Abstract**
This paper studies how people infer a state of the world from information structures that include additional, payoff-irrelevant states. For example, learning from a customer review about a product's quality requires accounting for the reviewer's otherwise-irrelevant taste. This creates an attribution problem common to all information structures with multiple causes. We report controlled experimental evidence for pervasive overinference about states that affect utility—a form of "omitted variable bias" in belief updating, providing an explanation for various misattribution patterns. In studying why systematic misattribution arises, we consistently find that errors are not due to deliberate effort avoidance or a lack of cognitive capacity. Instead, people behave as if they form incomplete mental models of the information structure and fail to notice the need to account for alternative causes. These mental models are not stable but context-dependent: Misattribution responds to a variety of attentional manipulations, but not to changes in the costs of inattention. (JEL: C91, D01, D83, D84)

## 1. Introduction

The difficulty of attending to, aggregating, and processing the abundance of available information in practice motivates a strand of work on errors in belief formation. For example, people may be partially inattentive to information (Sims 2003; Malmendier and Lee 2011; Hanna, Mullainathan, and Schwartzstein 2014; Caplin and Dean 2015; Bartoš et al. 2016; Enke 2020) or fail to account for the relationship between different signals (Eyster and Rabin 2010; Levy and Razin 2015; Enke and Zimmermann 2019). In some situations, however, the amount of information is manageable in principle, and agents are capable of attending to all available pieces of information. Rather than selecting or aggregating these signals, the challenge of belief formation then often

*The editor in charge of this paper was Nicola Pavoni.*

Acknowledgments: Financial support from the Sloan Foundation and the Institute on Behavior and Inequality (briq) is gratefully acknowledged.

E-mail: tgraeber@hbs.edu

*Journal of the European Economic Association* 2022 00(0):1–33
https://doi.org/10.1093/jeea/jvac052

lies in selecting the right interpretation of a piece of information. In this case, the agent faces an *attribution* problem as he may struggle to figure out what a given piece of information actually means. Rather than the first type of environment with *many signals*, this paper studies attribution problems in information structures with *many causes* for a single signal. To take a stylized example, suppose that a shopper reads a positive customer review that is a function of actual product quality and the reviewer's personal taste. Learning from a positive review about underlying quality requires accounting for other extraneous causes in the information structure, such as differing tastes. A failure to account for alternative causes creates misattribution to the causes of interest, a form of "omitted variable bias" in belief formation. For example, a decision-maker who does not factor in the role of varying tastes over-attributes a positive review to high product quality.

A collection of separately documented empirical findings is suggestive of this type of error. The defining pattern is excessive inference about a specific cause of interest while neglecting alternative causes that are "nuisance" from the decision-maker's perspective. For example, CEOs and politicians are rewarded for luck because performance evaluations and voter support partly fail to condition on external conditions such as the business climate (Bertrand and Mullainathan 2001; Wolfers 2002). People overstate the role of intentions relative to contextual factors and chance when explaining the behavior of others (Ross 1977; Gurdal, Miller, and Rustichini 2013), known as the fundamental attribution error in psychology. Applied work on attention shows that people often underreact to certain elements of the price structure, such as sales taxes, when learning from a price about the value of a good (Chetty, Looney, and Kroft 2009; Abaluck and Gruber 2011; Allcott 2011; Taubinsky and Rees-Jones 2018). When explaining the world, we tend to narrowly focus on the determinants that matter most to us, which may result in us attributing excessive causal power to them.

This paper tackles two questions. First, how do people learn about a target state of the world from information that also depends on otherwise irrelevant states? In simple, tightly controlled updating experiments, we document a systematic neglect of nuisance causes and misattribution to causes of interest. We validate the generalizability of the key finding using a naturalistic variant of the experiments that exploits an economically relevant situation and does not rely on explicit computations. Second, why does such misattribution arise? We examine this question by leveraging the distinction between "frictions" and "mental gaps" (e.g. Handel and Schwartzstein 2018).[1] Frictions are directly linked to the costs of information processing, and may occur due to mental processing noise, capacity constraints, or other forms of attentional limitations (Sims 2003; Gabaix 2014; Caplin and Dean 2015; Matějka and McKay 2015; Woodford 2019). A mental gap describes the divergence between how people think about a problem and how they should think about it given costs. People sometimes appear to

1. This taxonomy is representative of a collection of related classifications put forward in the literature, such as that of "bounds errors" versus "astray errors" (Rabin 2013).

form incorrect mental models or problem representations.[2] The data from more than 20 experimental treatments that examine the cognitive mechanisms underlying the neglect of alternative causes consistently point to a mental gap: Subjects are unaware of their neglect and minor attentional manipulations successfully debias respondents, whereas variations of the costs and benefits of attention have little to no effect.

We present causal evidence from laboratory and online experiments. We use a two-pronged approach with two complementary paradigms: The baseline experiment strips away the real-world context and associated ambiguities to create a maximally controlled belief updating setting, which comes at the cost of a potential lack of realism. The complementary set of vignette experiments replicates the results in settings with naturalistic task framing that is closer to real-world inference problems but sacrifices some of the control obtained in the former. In the baseline condition of the laboratory experiment, treatment *Narrow*, subjects guess an unknown, random state of the world and are paid for accuracy. Before indicating their guess, they receive a piece of information (the *signal*) that depends on both that target state and another unobserved state. Specifically, two numbers $X$ and $Y$ are drawn from known distributions. In this baseline condition, subjects have to guess $X$, but not $Y$. Because $Y$ is not a prediction target, it constitutes a *nuisance variable* from the subject's perspective: It confounds information about the target variable $X$, but its realization does not affect his payoff given a stated belief.[3] In a typical task, $X$ is drawn from the simple discretized uniform distribution on $\{30, 40, 50, 60, 70\}$ and $Y$ is drawn at random from $\{10, 20, 30, 40, 50, 60, 70, 80, 90\}$. Subjects observe a signal that depends on both states, such as the average of the drawn numbers, $S = (X + Y)/(2) = 70$. Crucially, inference from $S$ about $X$ requires accounting for the random variation in $S$ that is due to $Y$. In the context of the previous example, a shopper might want to infer unobservable product quality ($X$) from an observable customer review ($S$), which is a function of both quality $X$ and the reviewer's tastes $Y$. Failing to properly account for the stochasticity of $Y$ generates misattribution of the signal to $X$. Subjects are informed of the simple data-generating process and the signal structure, eliminating all structural uncertainty in the information environment. We confirm that subjects are not confused about the task setup using an extensive set of control questions; we always show all relevant information on the decision screen; and we run additional control treatments to address potential misunderstandings. In this baseline condition, where subjects are incentivized to state the full distribution of their beliefs about $X$ but not about $Y$, beliefs about $X$ exhibit pervasive neglect of the nuisance variable $Y$. In the numerical example above, this is equivalent to stating that $X = 70$ with certainty, as if $S = X$. The Bayesian posterior belief about $X$, by contrast, assigns equal probability to 50, 60, and 70. Across all tasks, only 17% of all stated beliefs are in line with

2. See, for example, Enke (2020), Jehiel (2005), Spiegler (2016, 2017), Gagnon-Bartsch, Rabin, and Schwartzstein (2019), Gennaioli and Shleifer (2010), Schwartzstein (2014), Eyster and Rabin (2010), Bordalo, Gennaioli, and Shleifer (2020b), and Hauser and Schwarz (2016).

3. We define a nuisance variable in Section 3.1.2 as one whose realization does not affect utility conditional on an action.

the Bayesian benchmark, whereas 62% display full neglect of $Y$. We refer to this as nuisance neglect and conceptualize its relationship to other forms of bias below.

In a baseline control treatment, *Broad*, a separate set of subjects is incentivized to guess the joint distribution of $X$ and $Y$, rather than only $X$. This turns $Y$ from a nuisance into a target variable, while keeping the overall monetary stakes as well as the objective updating problem (and thus the Bayesian posterior) exactly identical to the baseline condition. Because the information structure is unchanged, the complexity and cost of computing a posterior for $X$ should be unchanged. Note that this treatment is a control condition that alleviates a shortcoming of recent experimental work on belief formation because it manages to hold the objective updating problem constant across conditions.[4] In treatment *Broad*, we document a large and statistically significant treatment difference relative to treatment *Narrow*. Moreover, the median belief in *Broad* is indistinguishable from the Bayesian posterior, implying that the experimental setup is not too complex per se and subjects are in principle able to solve the task correctly. More than 70% of all stated beliefs in *Broad* correspond to the Bayesian posterior.

We examine the external validity of these findings using a set of naturalistic vignette experiments that leverage real-world scenarios, do not have the character of a math problem, include an application with economic relevance and a variant featuring a simple choice instead of a belief incentivized with a complex scoring rule. Next to the vignette experiments, a battery of laboratory and online experiments (i) tests the robustness of nuisance neglect by varying various elements of the experimental design, such as the specific signal structure (e.g. a signal outside of the variables support), the distribution of the random states (non-uniform distributions), and the elicitation procedure[5]; (ii) documents nuisance neglect in a large and heterogeneous online population; and (iii) tests the predictions of existing theories of belief formation in this setup. Specifically, we design sharp tests of different models using systematic variations of the data and signal structures. We find that the pattern of neglect of nuisance variables in the data is not consistent with overweighting the signal (Benjamin 2019 for a review of overinference), underweighting the base rate (Bar-Hillel 1980; Grether 1980), or diagnosticity-based theories of expectation formation (Bordalo, Gennaioli, and Shleifer 2018).

The second part of this paper studies *why* nuisance neglect arises by investigating the underlying cognitive mechanisms. We adopt the distinction between frictions and mental gaps as an instructive taxonomy for the present application: The neglect of nuisance variables may be due to the (computational) difficulty of accounting for the nuisance variable $Y$ in conjunction with $X$—a friction—or due to a failure to recognize the necessity to take into account $Y$ to begin with—a form of misconstrual or mental gap.

4. Experimental work on belief updating routinely compares beliefs in information environments with and without a feature of interest. A manipulation of the signal structure, however, can confound the analysis if it affects other properties of the updating problem, such as the complexity of forming an update (as in e.g. Enke and Zimmermann 2019).

5. For example, we disentangle the elicitation procedure from prediction incentives by having subjects state the joint distribution when only $X$ is incentivized—unlike in *Narrow*, or by having them state a marginal belief about $X$ first when both variables are incentivized—unlike in *Broad*.

To examine these explanations, we design additional experiments that test for a potential mental gap. We present a series of additional experiments that aim to manipulate how people think about the updating task while keeping the cost of accounting for $Y$ constant. If drawing people's attention to the role of $Y$ without changing the updating problem affects the degree of nuisance neglect, then a mental gap is likely to play a role.

We present three main findings from the analysis of mental gaps. First, nuisance neglect is reduced substantially and beliefs are pre-dominantly Bayesian once attention is drawn to $Y$. A contextual cue to attend to $Y$ is sufficient to reduce nuisance neglect while maintaining $Y$'s role as a nuisance variable and holding constant the difficulty of accounting for it. In treatment *Hint*, subjects only guess $X$ but see an additional verbal statement on each elicitation screen: "Also think about the role of $Y$". The hint produces a large and statistically significant treatment difference relative to the baseline condition *Narrow*.

Second, while the exogenous manipulations of attention have the potential to debias, we find that subjects are able to overcome nuisance neglect on their own when nudged to reconsider their solution strategy. In treatment *Enforced Deliberation*, we implement a 30-second deliberation time on the elicitation screen before the input fields are activated. The objective is to encourage subjects to deliberate their problem interpretation *before* they form their posterior. Enforced deliberation time substantially reduces nuisance neglect and is roughly half as effective as an explicit hint.

Note that the effect of minor attentional manipulations is striking in the sense that even in condition *Narrow*, all relevant pieces of information are displayed on the screen and we ensure that subjects are not confused by the setup. However, they may still fail to realize the necessity to account for the variation of $Y$ to begin with and would consequently be *unaware* of committing an error. In a third step, we test this lack-of-awareness hypothesis directly by measuring confidence in beliefs using incentivized willingness-to-pay (WTP) to have a guess replaced by an optimal guess. Exploiting causal variation, we find that nuisance neglect is associated with similar confidence levels as Bayesian updating, indicating that subjects are unaware of the neglect.

These three findings consistently suggest an underlying mental gap: Attentional manipulations that plausibly hold the cost of information processing fixed close the mental gap of failing to attend to $Y$, which subjects seem to be unaware of to begin with.

In a companion exercise, we empirically investigate the friction mechanism for nuisance neglect. Why do people systematically neglect elements of an information structure, even in simple contexts? One candidate explanation is that such model simplifications reflect a strategy to economize on cognitive costs.[6] We report two main findings on the applicability of this cognitive cost–benefit perspective to our setting.

6. A prominent view in cognition research holds that humans are "cognitive misers" who continuously seek strategies to avoid thinking (Fiske and Taylor 2013). Similarly, a large class of models in economics relies on weighing the expected benefits against the cognitive costs of attention (Gabaix 2014; Caplin and Dean 2015), prominently including theories of rational inattention (Sims 2003, 2006).

First, we find that increasing the stake size ten-fold (in the laboratory) or five-fold (in online experiments) substantially increases effort as measured by response times but does *not* reduce the prevalence of nuisance neglect, at odds with an underlying lack of effort. Second, we directly test for the presence of cost–benefit considerations by manipulating the specific monetary loss incurred from committing nuisance neglect, based on how much "noise" and "bias" the presence of $Y$ introduces into the posterior of an agent who mistakenly updates as if $S = X$. Strikingly, the presence of nuisance neglect does *not* respond to the monetary loss associated with its expected (in)accuracy.

Taken together, the analysis of mechanisms suggests that nuisance neglect occurs when subjects do not mentally account for $Y$ to begin with. They are unaware of this omission, and it does not reflect a lack of effort. People seem to initially "fail to notice" the necessity of accounting for the variation in $Y$, which may lead them to form a misspecified problem representation. Attentional cues that nudge subjects into re-considering the problem (conditions *Enforced Deliberation* and *Hint*) improve updating substantially. The combined evidence is more consistent with a mental gap interpretation of misattribution.

This paper proceeds as follows. Section 2 embeds this paper in the existing literature. In Section 3, we present the baseline design and results from the laboratory and online experiments, as well as extensions that include a replication in a naturalistic context and robustness exercises. In Section 4, we examine why nuisance neglect occurs based on the distinction between mental gaps and cost–benefit considerations. Section 5 concludes.

## 2. Related Literature

This paper contributes to several literatures. In the experimental literature, this study of misattribution in the basic case of interpreting a *single* piece of information complements recent work on updating errors in situations that require the joint processing and aggregation of *many* pieces of information (Enke and Zimmermann 2019; Enke 2020). Other related work highlights failures of hypothetical thinking (Martínez-Marquina, Niederle, and Vespa 2019; Esponda and Vespa 2014, 2019) and the failure to notice important features of the available data (Hanna, Mullainathan, and Schwartzstein 2014). Benjamin (2019) reviews a voluminous body of empirical research on probabilistic reasoning. His meta-study concludes that beliefs often tend to be less sensitive to variation in problem parameters—such as the base rate, diagnosticity, and sample size—than postulated by Bayes' rule. That people do respond to parameters albeit too little differs from the type of discrete neglect of a part of the signal structure documented here. Moreover, we show that subjects do not follow a compelling intuition when committing nuisance neglect, which underlies many judgement errors studied in the heuristics and biases literature (Kahneman and Tversky 1982; Tversky and Kahneman 1983; Morewedge and Kahneman 2010). Finally, this paper contributes a new perspective to the long-standing debate on the conditions for overreaction versus underreaction to information (Coibion and Gorodnichenko 2012;

Greenwood and Shleifer 2014; Coibion and Gorodnichenko 2015; Frydman and Nave 2016; Adam, Marcet, and Beutel 2017; Landier, Ma, and Thesmar 2017; Bordalo et al. 2020a).[7] Nuisance neglect simultaneously generates overreaction to payoff-relevant causes and underreaction to nuisance causes, providing testable predictions on their relative likelihood of occurrence.

In studying *why* updating errors occur, our findings on the source of nuisance neglect in attribution problems square with those of Enke (2020), who finds that people sometimes narrowly focus on visible parts of the information structure in signal aggregation tasks. Enke (2020) argues that people form simplified mental models of a problem that respond to the computational complexity of a task. Comparable findings in the updating environments studied here hint at a common cognitive mechanism underlying belief errors in both signal aggregation and attribution problems: an unwitting neglect of parts of the structure of updating problems. While Enke (2020) shows that this neglect can be context-driven by varying which signals are visible, the present paper highlights a different channel: Heuristic model simplifications may often be determined by people's incentive structure, irrespective of which parts of the information structure are visible.

On the applied side, this paper speaks to a collection of separately documented misattribution patterns. One line of work starting with Chetty, Looney, and Kroft (2009) shows inattention to specific features of the decision context (Abaluck and Gruber 2011; Allcott 2011; Abaluck and Adams 2017; Taubinsky and Rees-Jones 2018). For example, Taubinsky and Rees-Jones (2018) find that people underreact to sales taxes. While their experiment does explicitly pose an inference problem, the results are consistent with consumers systematically overinferring from price ($S$) about the value of the product ($X$) while neglecting the sales tax ($Y$). Other phenomena that can be interpreted through the lens of nuisance neglect are: outcome bias in punishing decisions that are based on luck ($Y$) rather than effort ($X$) alone (Gurdal, Miller, and Rustichini 2013; Brownback and Kuhn 2019); in consumer choice, there is misattribution of positive experience to the intrinsic value of an outcome ($X$) while neglecting reference-dependent surprise ($Y$) (Bushong and Gagnon-Bartsch 2022) and to the quality of a consumption good ($X$) rather than a contextual state such as the weather ($Y$) (Haggag et al. 2018); and in social learning contexts, people overinfer about a person's private information ($X$) from their action, neglecting that the action also embeds private information from earlier movers ($Y$) (Eyster, Rabin, and Weizsäcker 2018). In much of this work, inattention specifically occurs to problem features that are plausibly nuisance variables. At the same time, findings from previous experimental work on environments with many signals do not apply to these settings (e.g. Bartoš et al. 2016; Enke and Zimmermann 2019).

7. Championed by Kahneman and Tversky and prominent in finance is the view that beliefs move too much (Tversky and Kahneman 1971; Shiller 1981; De Bondt and Thaler 1985; Bordalo et al. 2019), while an older psychology literature and the dominant view in macroeconomics maintains that beliefs tend to move too little (Edwards 1968; Rabin and Schrag 1999; Mankiw and Reis 2002; Benjamin 2019).

Research in cognitive science has studied related phenomena that speak to the external validity of our results. The "causal frame problem" shows that people often form incomplete causal models of a problem. Work on biases in causal reasoning finds that people employ cognitive shortcuts that can result in the neglect of alternative causes (Fernbach, Darlow, and Sloman 2010; Sloman and Lagnado 2015; Fernbach and Rehder 2013). This body of work indicates that the findings from the highly controlled but stylized experimental environments studied here carry over to environments with a more naturalistic task framing.

Finally, this paper speaks to a large theoretical literature. Work on the rational inattention paradigm focuses on rational information acquisition given cognitive capacity constraints or processing costs. Rational inattention models do not generate systematic misinference conditional on processing a piece of information as they posit Bayesian inference from the information that an agent actually attends to (Sims 2003; Wiederholt 2010; Matejka and McKay 2014; Caplin and Martin 2015; Caplin et al. 2020). The discrete neglect of certain dimensions in the data is reminiscent of the sparsity-based model of Gabaix (2014), applied to belief updating. The lack of responsiveness to variations in costs and benefits, however, is at odds with sparse maximization. A key characteristic of the combined evidence is that people appear to form inaccurate mental representations of problems because they are looking at the problem the wrong way, rather than trading off the benefits and costs of more accurate representations. This appears more compatible with frameworks of mental gaps than with models of rational inattention. Inaccurate priors may lead to self-serving misattribution (Hestermann and Yaouanq 2020) or discrimination (Chauvin 2020) but are unlikely to be at play here because priors are controlled experimentally. The most closely related theoretical frameworks view incomplete representations as reflecting incorrect beliefs about which variables matter (Schwartzstein 2014; Gagnon-Bartsch, Rabin, and Schwartzstein 2019). All of these models share the prediction that representations should look fairly consistent across problems. The evidence in this paper highlights that they miss how heuristic model simplifications may often not be stable but constructed *on-the-fly* in response to task demands, environmental cues, and even suggestions to reconsider a representation.

## 3. Evidence for Nuisance Neglect

### *3.1. Baseline Experiments*

To causally examine the role of nuisance variables in information structures for belief updating, the experimental design aims to satisfy the following requirements: (i) a fully controlled and transparent data-generating process and information structure that is known to subjects, (ii) an experimental manipulation of the presence of nuisance causes, (iii) limited complexity to minimize confusion, and (iv) an incentive-compatible procedure to extract beliefs.

TABLE 1. Overview of baseline task specifications.

| Sample space of $X$ | Sample space of $Y$ | Signal structure | Signal realization |
|---|---|---|---|
| 30, 40, 50, 60, 70 | 10, 20, 30, 40, 50, 60, 70, 80, 90 | $(X + Y) \div 2$ | 60 |
| 230, 240, 250, 260, 270 | 210, 220, 230, 240, 250, 260, 270, 280, 290 | $(X + Y) \div 2$ | 230 |
| 180, 190, 200, 210, 220 | 180, 190, 200, 210, 220 | $(X + Y) \div 2$ | 200 |
| 80, 90, 100, 110, 120 | −30, −20, −10, 0, 10, 20, 30 | $X + Y$ | 80 |
| 130, 140, 150, 160, 170 | −25, −15, −5, 0, 5, 15, 25 | $X + Y$ | 165 |

Notes: Overview of the five baseline belief tasks in the laboratory study. The distributions of $X$ and $Y$, as well as the signal structure, are identical in both treatment conditions. $X$ and $Y$ are independently drawn from two discrete uniform distributions; that is, every indicated outcome is equally likely. In the baseline study, all subjects received the same (random) signal realization. In the complementary online experiments, signal realizations were drawn at the subject level.

*3.1.1. Design.* Experimental variation in the presence of nuisance causes can be achieved by changing the information structure, but this also affects the complexity of updating beliefs across conditions. We instead design a simple setting that implements this variation without changing the information structure or data-generating process. The basic updating task features two unobserved random numbers, $X$ and $Y$, generated by stochastic processes known to subjects. To simplify, these numbers are independently drawn from two discrete uniform distributions with small sample spaces. Subjects receive a signal $S = s$ on the two unknown draws: Depending on the task, they see either the sum or the average of the two numbers.[8] The signal structure maps two inputs, that is, the realizations of random variables $X$ and $Y$, to a one-dimensional output, that is, the observed signal $s$. The experiment creates exogenous between-subject variation in whether the agent's payoff depends on the realizations of only one or both of the inputs in the signal structure. There are two experimental conditions: In *Narrow*, subjects are paid to guess only $X$, while in *Broad*, subjects are paid to guess both $X$ and $Y$. The experimentally controlled prior, the signal structure, and the Bayesian posterior are identical in *Narrow* and *Broad*. A Bayesian agent thus forms identical beliefs in both conditions. By randomly choosing only one of the guesses in *Broad* for payment, the size of the monetary incentive is kept constant.

Subjects complete the five updating tasks of Table 1 in random order without receiving feedback in between. For example, in the first task of Table 1, $X$ is one of

8. In the baseline tasks, the signal is an unbiased estimator of the mean of $X$. Either subjects receive the average of the drawn numbers and the prior distributions of $X$ and $Y$ have identical means, or they see the sum of the drawn numbers and $Y$ has a mean of zero. We study more general signal structures from Section 3.2.

five numbers: 30, 40, 50, 60, or 70 with equal probability, while $Y$ is independently drawn with equal probability from 10, 20, 30, 40, 50, 60, 70, 80, and 90. Subjects learn that the average of $X$ and $Y$ is 60 and then state their belief as described in detail in Section 3.1.3.[9] To solve this task, subjects need to identify all $(X, Y)$ combinations with an average of 60, that is, (30, 90), (40, 80), (50, 70), (60, 60), and (70, 50). Both numbers being drawn uniformly and independently, it follows that each of these outcomes is equally probable. Additional task specifications and treatment variations address the robustness of the baseline results and examine the nature of updating rules (see Section 3.2). For example, we replicate our findings with more general data and signal structures, for example when the signal falls outside of the support of a variable.

A key feature of this design is that unlike related empirical studies of updating errors (Caplin, Dean, and Martin 2011; Dean and Neligh 2019; Enke and Zimmermann 2019; Enke 2020), this experimental setup holds the information structure fixed across conditions.

*3.1.2. Definition of Nuisance Variables and Predictions.* To fix ideas, we delineate basic concepts underlying the baseline treatment comparison in a setting that loosely follows Gabaix (2019). Assume an agent (he) who states a belief $b \in \mathbb{R}$ about the two-dimensional vector of random states in the updating task, $(x, y) \in \mathbb{R}^2$. Without loss of generality, we normalize $\mu_X = \mu_Y = 0$, such that $(x, y)$ denote deviations from their respective means. The agent chooses $b$ to maximize linear-quadratic utility after observing a signal $s$ about the random draw $(x, y)$. Crucially, the signal is generated by a deterministic function of *both* random variables, that is, $s = f(x, y)$, with $\frac{\partial s}{\partial x} \neq 0$ and $\frac{\partial s}{\partial y} \neq 0$. The utility function

$$u(b, x, y) = -\frac{1}{2}\left(b - \eta_x x - \eta_y y\right)^2 \tag{1}$$

yields the following optimal belief:

$$b^r(s) = \max_b \; \mathrm{EU}|s = \max_b \; \mathbb{E}\left[-\frac{1}{2}\left(b - \eta_x x - \eta_y y\right)^2 |s\right] \tag{2}$$

$$= \mathbb{E}[\eta_x x + \eta_y y|s] = \eta_x \mathbb{E}[x|s] + \eta_y \mathbb{E}[y|s]. \tag{3}$$

This optimal belief is a function of the Bayesian conditional posterior expectation of $x$, $\mathbb{E}[x|s]$, and $y$, $\mathbb{E}[y|s]$, as well as weight parameters $(\eta_x, \eta_y)$ that reflect how strongly the agent's utility depends on the realization of each variable. The definition of a nuisance variable directly follows from the weight parameters.

DEFINITION. *$Z \in \{X, Y\}$ is a nuisance variable in an updating problem if its realization $z$ does not affect the agent's expected utility conditional on a stated belief. Formally, $\forall b \in \mathbb{R}$: $\frac{\partial[\mathrm{EU}|b]}{\partial z} = 0$. This is the case iff $\eta_z = 0$.*

9. Note that in the baseline study, all subjects received the same (random) signal realization. In complementary online experiments, signal realizations were drawn at the subject level.

Intuitively, the agent's expected payoff in a belief formation task does not respond to the realization of a nuisance variable for any given stated belief. A nuisance variable is payoff-irrelevant after stating a belief. Importantly, this does not mean that a nuisance variable is irrelevant for the agent's *optimal* belief $b^r$, which is clear from equation (3): Even if $\mu_y = 0$, $b^r$ mechanically depends on $y$ through the conditional expectation $\mathbb{E}[x|s]$. $\mathbb{E}[x|s]$ is determined by the signal structure $S$ that we assumed is a function of $Y$.[10] The definition of a nuisance variable highlights that optimal beliefs can depend on variables whose realizations are payoff-irrelevant conditional on a stated belief. This points to a crucial distinction between incentives provided through payoffs on the one hand, and the necessity of taking into account all elements of an information structure to form a Bayesian posterior on the other hand.

We now apply this idea to the treatment variation. In *Broad*, the agent is paid for the accuracy of his joint posterior about $(x, y)$, so that his utility depends on both realizations given a stated belief, that is, $\eta_x \neq 0$ and $\eta_y \neq 0$. Thus, neither $X$ nor $Y$ are nuisance variables. In *Narrow*, however, the agent's expected utility given a belief only depends on the realized state of $X$ but not on that of $Y$, that is, $\eta_x \neq 0$ but $\eta_y = 0$.

Note, however, that since the Bayesian belief about $(x, y)$ is independent of the prediction incentives, the treatment manipulation is designed in such a way that it is inconsequential under Bayesian updating. While we focus on a tightly controlled, stylized setting for the reasons outlined above, note that nuisance variables are readily identifiable in applied contexts: They are sources of stochasticity that are materially irrelevant to an agent beyond the necessity to account for them in an inference problem.

The thrust of the baseline prediction is that people neglect nuisance variables in the updating problem. A priori, this neglect of $Y$ could take on a number of different forms. The decision-maker may implicitly neglect the variance of $Y$, replace $Y$ with a "default" value (as in Gabaix 2014), or apply a particular non-Bayesian updating rule. We investigate the precise form of nuisance neglect using additional experimental variations, see Section 3.2.1. A candidate form of neglect is that the agent interprets the signal as if it only depends on $X$, but not $Y$. Nuisance neglect in condition *Narrow* may then be characterized by the agent taking the signal as fully revealing about $X$, *as if* generated by an alternative deterministic signal structure $\widetilde{S} = g(X)$. The neglectful agent forms his belief based on a flawed posterior $P(X|\widetilde{S} = s)$ instead of $P(X|S = s)$.

PREDICTION 1. *Beliefs exhibit nuisance neglect.*

(a) *Beliefs in condition Narrow imply a neglect of $Y$. Specifically, subjects take the signal as fully revealing about $X$.*
(b) *Beliefs in condition Broad are Bayesian.*

Prediction 1 directly implies a treatment difference between stated beliefs in conditions *Narrow* and *Broad*. The above simplistic notion of neglecting states that are

10. Note that $\mathbb{E}[x|s] = \mathbb{E}[x \mid f(X, Y) = s]$.

payoff-irrelevant given a stated belief abstracts from the specific features of attention. It merely serves to set the stage for our in-depth analysis of the nature of attention and the more explicit framework of the origins of neglect. In particular, we will later argue and show that being a nuisance variable is not a sufficient condition for neglect in the inference problem and disentangle between endogenously chosen attention and exogenous attentional cues.

*3.1.3. Procedures.* Subjects in condition *Broad* guess the joint distribution of $X$ and $Y$ and are randomly paid for their accuracy in guessing either of these (decision screen in Online Appendix Figure G.4). Subjects in condition *Narrow* only guess the marginal distribution of $X$ (Online Appendix Figure G.1).[11] The design unobtrusively obfuscates the study's objective: Subjects receive their signal in encrypted form and have to decipher it using a simple two-step decoding protocol.[12] Note that no subject failed to implement the protocol. In a control treatment (*Simplification*, see also Online Appendix C.4) and all online experiments (Section 3.2.1), this feature was removed. The findings absent the obfuscation indicate that the obfuscation would not have been necessary in the baseline experiment. Each belief elicitation (excluding the deciphering stage) is subject to a five-minute time limit. The findings are robust to removing both the deciphering and the time limit (Section 3.2).

The elicitation procedure aims at providing a full characterization of subjective beliefs by having subjects indicate the entire posterior distribution instead of a point prediction. At the end, one of the tasks is randomly selected to be paid out based on the Binarized Scoring Rule with a prize of 10 euros (Hossain and Okui 2013).[13] Subjects receive extensive instructions and have to complete eight control questions that test their understanding of the instructions, the data-generating process and signal structure, as well as the elicitation protocol (see Online Appendix G). In two unpaid practice tasks, subjects are trained to indicate a verbally described belief in a way that maximizes their payoff. This training stage is identical across treatments.

The belief updating problems are followed by a questionnaire. To shed light on correlates of subject-level heterogeneity in belief formation, we measure performance on an incentivized test of cognitive capacity (10 Raven matrices, 0.2 euros per correct answer) and elicit a measure of risk preferences (Falk et al. 2016).

11. There is a treatment difference in the elicitation protocol, that is, whether $X$ and $Y$ or only $X$ is elicited. Additional treatment variations harmonize the elicitation protocol, that is, subjects with *Narrow* incentives predict both $X$ and $Y$, and subjects with *Broad* incentives predict first the marginal of $X$, and then the marginal of $Y$ on a separate subsequent page. All main findings persist. See Section 3.2.

12. Subjects see a sequence of letters. First, each letter has to be translated into a digit based on a decoding key displayed on the screen. Then the number 20 has to be added to the result. Subjects are familiarized with the deciphering process in the practice stage. See also the instructions in Online Appendix G.

13. The scoring rule proposed by Hossain and Okui (2013) remains incentive compatible if subjects are risk averse. We adopt the approach suggested by Hossain and Okui (2013) to incentivize the entire stated distribution based on the sum of squared deviations between the probability mass allocated to each value of the distribution and the corresponding mass that should be allocated after learning the realized outcome. See Online Appendix G for further details.

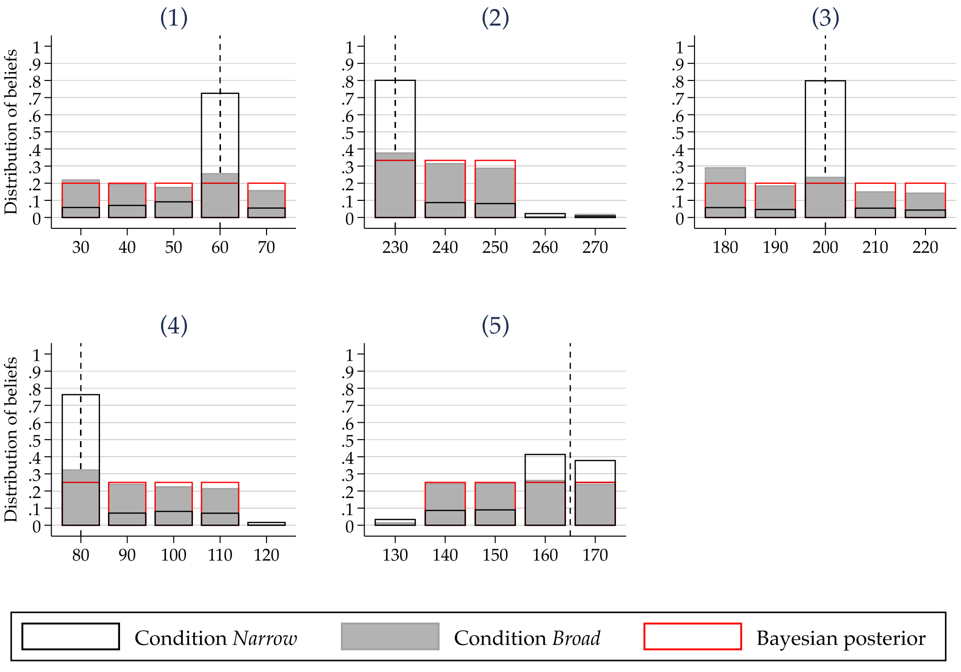


FIGURE 1. Treatment averages of stated belief distributions about $X$ for each of five baseline tasks. $N = 72$ for each treatment in each task. The horizontal axis shows possible outcomes of $X$. The Bayesian posterior belief is provided for reference. The observed signal is indicated by the vertical dashed line. In all five tasks, $X$ and $Y$ follow independent discrete uniform distributions. The task order was randomized at the subject level. The distributions and signal structure for each task are shown in Table 1. Subjects observed the mean of the drawn numbers in tasks (1), (2), and (3), and they saw the sum in (4) and (5).

A total of144 student subjects (72 in each treatment) participated in six sessions of the baseline experiment run at the University of Bonn's BonnEconLab in July 2017. Treatment status was randomized within session. We implemented the study in oTree (Chen, Schonger, and Wickens 2016). Mean earnings amounted to 11.40 euros—including a 5-euro show-up fee—for an average session duration of 57 minutes.

*3.1.4. Results.* We begin with an analysis of stated beliefs at the aggregate level before exploring their heterogeneity in Section 3.1.5. Figure 1 illustrates raw beliefs in each baseline task. It shows the sample average of stated belief distributions in both treatment conditions, alongside the Bayesian belief and the signal realization. The average subject in *Broad* forms beliefs that are closely aligned with the Bayesian posterior. In *Narrow*, by contrast, subjects on average assign too much probability mass to outcomes close to the signal value, as implied by inattention to $Y$.

Table 2 provides summary statistics and non-parametric tests by task. Median beliefs in *Narrow* (column (3)) and *Broad* (column (4)) closely correspond to the

TABLE 2. Beliefs about $X$ in baseline tasks.

| Signal realization | Bayesian posterior distribution | Stated posterior distribution | | Sign test of median | | Mann–Whitney $U$ test |
|---|---|---|---|---|---|---|
| | | *Narrow* $N = 72$ | *Broad* $N = 72$ | *Narrow* vs. Bayesian | *Broad* vs. Bayesian | *Narrow* vs. *Broad* |
| (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| | *Distribution mean (distribution variance)* | *Median of distribution means (median of distribution variances)* | | *p-value: distribution of means (p-value: distribution of variances)* | | |
| 60 | 50 | 60 | 50 | <0.001 | 0.664 | <0.001 |
| | (200) | (0) | (200) | (<0.001) | (0.011) | (<0.001) |
| 230 | 237.6 | 230 | 240 | <0.001 | <0.001 | <0.001 |
| | (71.7) | (0) | (67) | (<0.001) | (<0.001) | (<0.001) |
| 200 | 200 | 200 | 200 | 1.000 | 0.012 | 0.024 |
| | (200) | (0) | (200) | (<0.001) | (0.004) | (<0.001) |
| 80 | 95 | 80 | 95 | <0.001 | 0.508 | <0.001 |
| | (125) | (0) | (125) | (<0.001) | (0.180) | (<0.001) |
| 165 | 155 | 165 | 155 | <0.001 | 1.000 | <0.001 |
| | (125) | (25) | (125) | (<0.001) | (0.180) | (<0.001) |

Notes: Beliefs in *Narrow* and *Broad* by task. Each stated belief is a distribution, summarized here by its mean and variance. The table shows medians of stated distribution means and stated distribution variances for each condition (columns (3) and (4)), and compares these to the mean and variance of the Bayesian posterior distribution (columns (5) and (6)). Column (7) shows treatment comparisons. The task order was randomized at the subject level.

observed signal realization (column (1)) and the Bayesian benchmark (column (2)), respectively. Column (7) shows that belief distribution means and belief distribution variances are significantly different between treatments at the 0.1% level (Mann–Whitney $U$ tests).[14] Note that the median variance of stated distributions in *Narrow* is far too low, indicating that subjects hold too precise beliefs.

RESULT 1. *Beliefs display nuisance neglect.*

(a) *The median belief in Narrow exhibits exact nuisance neglect, that is,* $P(X|X = s)$.

(b) *The median belief in Broad equals the Bayesian posterior.*

(c) *There are significant treatment differences in stated posterior distributions between Narrow and Broad.*

Three implications of these results are that (i) there is no systematic confusion about the experimental setup, since the average belief in *Broad* is nearly Bayesian; (ii) in *Narrow*, the average belief *overshoots* in the direction of the signal; and (iii) is *overprecise* relative to both the Bayesian benchmark and beliefs stated in *Broad*. Overprecision is the common finding in belief research that the implied variance of stated beliefs is too low, indicating people's excessive confidence in their own judgments (Moore, Tenney, and Haran 2015). In the present context, overprecision in *Narrow* is solely generated by the presence of a nuisance variable, as the information structure does not change relative to *Broad*. Task (3) in Figure 1 exemplifies the role of overprecision. Since the signal realization coincides with the mean of the Bayesian posterior distribution, subjects in *Narrow* form unbiased beliefs on average about $X$, that is, they correctly guess the expected value of $X$ given the signal. However, they express too much certainty that this expected value of $X$ equals the actual draw. This finding could not be identified from point predictions about $X$ alone.

Nuisance neglect implies a sizeable monetary cost for subjects. The average expected payoff for the beliefs stated in the baseline tasks is 53% higher in *Broad* than in *Narrow* (5.86 vs. 3.82 euros, $p < 0.001$, Mann–Whitney $U$ test).[15]

*3.1.5. Heterogeneity.* Next, we examine what are *typical* beliefs in each condition. We characterize each stated belief by its relative proximity to the Bayesian posterior as opposed to the nuisance neglect posterior. Using that each observation is a distribution, we calculate the Hellinger metric distance (Hellinger 1909) between the stated posterior

14. This holds for all tasks except the distribution means in task (3), in which the signal realization coincides with the Bayesian posterior mean.

15. Actual earnings for the baseline tasks also significantly differ across groups (means of 4.56 in *Narrow* and 2.22 euros in *Broad*, $p = 0.005$, Mann–Whitney $U$ test), but these further depend on randomness induced by the binarized scoring rule as well as an additional choice by subjects that affects their payoff (see Section 4.2.3).

$b_X$ and the Bayesian posterior $P(X|S)$ distributions[16]:

$$H_B = \frac{1}{\sqrt{2}} \sqrt{\sum_{i=1}^{k} \left( \sqrt{b_{X_i}} - \sqrt{P(X_i|S=s)} \right)^2}. \tag{4}$$

Given an analogous distance to the inattentive posterior distribution, $H_N$,[17] we define an inattention score $\theta$ that captures the distance of the subjective belief distribution to the Bayesian distribution, relative to the sum of the distances of the subjective distribution to the inattentive and the Bayesian posterior:

$$\theta = \frac{H_B}{H_B + H_N}. \tag{6}$$

A Bayesian belief corresponds to $\theta = 0$ and nuisance neglect to $\theta = 1$. The parameter $\theta$ is computed individually for each stated belief. First, Figure 2 provides a histogram of empirical inattention parameters by treatment condition. This analysis pools all stated beliefs in a treatment condition across tasks and subjects. More than 70% of beliefs in *Broad* but less 20% in *Narrow* can be characterized as close to Bayesian ($\theta < 0.1$). By contrast, about 60% of beliefs in *Narrow* are close to nuisance neglect ($\theta > 0.9$), with the remaining 20% located in between the two extremes. The vast majority of stated beliefs are either fully sophisticated or fully inattentive to $Y$. This measure of inattention suggests a markedly bi-modal distribution of beliefs. Second, we analyze the within-subject heterogeneity of beliefs by counting how often each subject states a belief that is close to Bayesian ($\theta < 0.1$) or nuisance neglect ($\theta > 0.9$), as opposed to a belief that corresponds to neither of the two ($\theta \in [0.1, 0.9]$). A total of 58% of subjects in *Broad* but only 6% in *Narrow* state *all* of their beliefs in line with Bayes' rule. 44% of subjects in *Narrow* (but none in *Broad*) exhibit nuisance neglect in all of their stated beliefs. Consequently, a share of 62% in *Broad* and 50% in *Narrow* switch at least once between the three updating modes specified here. Kernel density estimates of the subject-level average of $\theta$ display a pronounced peak around zero mean inattention in *Broad* and a less pronounced clustering of subjects with mean inattention above 0.8 in *Narrow* (see Online Appendix Figure B.2). This suggests that while a considerable fraction of subjects are consistently inattentive, most subjects in *Narrow* exhibit some heterogeneity, with 15.5% reporting both a fully Bayesian belief and a belief implying exact nuisance neglect at least once. Strikingly, we find that a staggering 93% of beliefs stated in rounds that followed a close-to-Bayesian belief ($\theta < 0.1$) are also close to Bayesian. This fraction was only slightly lower in condition

16. The Hellinger distance is a bounded metric used to characterize the similarity between two probability distributions (Bandyopadhyay, Brittan, and Taper 2016). It is suited for the present purpose as it is a proper metric, unlike, for example, the Kullback–Leibler divergence, which does not satisfy symmetry.

17. $H_N$ is calculated as

$$H_N = \frac{1}{\sqrt{2}} \sqrt{\sum_{i=1}^{k} \left( \sqrt{b_{X_i}} - \sqrt{P(X_i|\widetilde{S}=s)} \right)^2}. \tag{5}$$

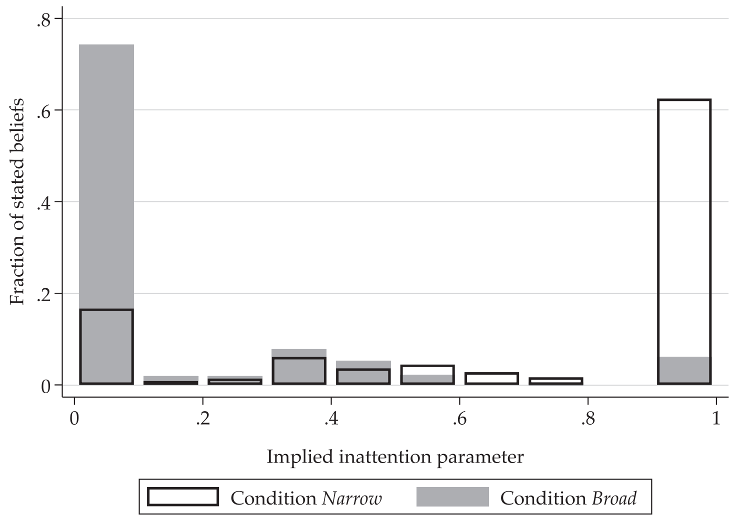


FIGURE 2. Distribution of implied inattention scores by treatment condition. $N = 1135$. Displayed are binned histograms for the implied inattention parameter based on all beliefs stated in the five baseline tasks. Inattention scores are calculated as $\theta = \frac{H_B}{H_B + H_N}$, where $H_B$ and $H_N$ denote the Hellinger distances of the stated belief distribution to the Bayesian posterior and the nuisance neglect posterior, respectively. A parameter of $\theta = 0$ corresponds to Bayesian updating. $\theta = 1$ implies nuisance neglect.

*Narrow* (82%) than in *Broad* (95%). This finding highlights the role of "insight": Once people figure out the right strategy, they consistently apply it throughout subsequent problems. This provides a first indication for the relevance of the mental model of the cognitive solution approach.

### *3.2. Robustness and External Validity*

The baseline study documents nuisance neglect in a specific configuration of the information environment and experimental setup. Using further experiments, we test the robustness of the findings, address potential confounds, and examine the generalizability of the baseline result.

*3.2.1. Robustness and Extensions.* This section summarizes a collection of robustness exercises and extensions that include (i) additional tasks introducing various departures from the simple discrete uniform case, (ii) a direct test of a signal anchoring heuristic, (iii) two treatments that exactly align the elicitation procedure across conditions, (iv) a simplified version that removes the deciphering stage and time limits, (v) a test of a face value heuristic, and (vi) an examination of the form of nuisance neglect across information structures. The following provides a brief discussion of these analyses, with all details relegated to the Online Appendix.

Adding to the baseline tasks in Table 1, four additional tasks were presented in random order after the baseline tasks. Highly significant treatment effects persist ($p < 0.001$, Mann–Whitney $U$ tests) under continuous uniform, normally distributed, or correlated data structures, or a case in which the signal realization is outside of the range of $X$. See Online Appendix C.1 for the robustness task specifications and detailed results.

A potential concern is that the treatment manipulation in the baseline study not only varies the incentive structure as postulated by the definition of a nuisance variable but also the elicitation procedure: Subjects in *Narrow* only state a belief about $X$, whereas subjects in $Broad$ guess both $X$ and $Y$.[18] To better understand the extent to which the treatment effect is due to the difference in elicitation procedures, two additional treatments are designed to obtain a 2 (incentives *Narrow* vs. *Broad*) $\times$ 2 (elicitation of only $X$ vs. $X$ and $Y$) between-subjects design. We find that given an incentive structure, that is, *Narrow* or *Broad*, harmonizing the elicitation protocol reduces the treatment effect by roughly one third, while all differences in estimated inattention scores remain highly significant (see Online Appendix C.3). Put differently, most of the treatment effect is driven by prediction incentives as opposed to the elicitation procedure.

Finally, drastically simplifying condition *Narrow* by removing the deciphering stage as well as all time limits induces a reduction in the implied inattention parameter ($p < 0.001$), but the treatment effect persists in a conservative comparison against the baseline condition *Broad*, which included both deciphering and time limits ($p < 0.001$; Online Appendix C.4).

A possible explanation of nuisance neglect is that subjects use the heuristic of reporting back the signal value, akin to exact anchoring or taking the signal at face value. Treatment *Computation* tests the face value explanation by adding a simple algebraic computation into the information structure, in such a way that it remains equally plausible to anchor on the observed signal value. For example, instead of $S = (X + Y)/(2)$, subjects receive the modified signal $S = (X + Y)/(2) - ((2 \times 10) + 30)$. We find minimal evidence for anchoring on the observed signal. Instead, subjects are able and willing to invert the computations, but still do not account for $Y$.[19]

Next, we complement the baseline evidence from the laboratory with online experiments in a large, more heterogeneous population. We implement several design modifications for these experiments that are discussed in Online Appendix D. While the design of the lab study was suitable for a sample of highly attentive student subjects, adjustments were necessary to adapt this experiment to the plausibly less attentive online worker population. Correspondingly, the online study did not serve the purpose

18. This is a deliberate design choice: making a prediction in itself provides a (non-monetary) incentive to pay attention, or constitutes a form of "cue" as studied in Section 4. Note that the information set at the time of stating a belief is held exactly constant across treatments, so that subjects in *Narrow* do not have to memorize the distribution of $Y$.

19. Further treatment details, figures, and results are relegated to Online Appendix C.2.

of an exact replication, but instead aimed at documenting the prevalence of nuisance neglect under less controlled conditions and a more diverse sample, and served as the basis for investigating different features of the phenomenon.

The main finding of substantial nuisance neglect replicates in the online study. Specifically, 53% of stated beliefs imply an attention parameter $\theta$ above 0.9. In addition, we document evidence for an additional updating mode, "signal neglect" or non-updating, a frequent finding in belief formation studies (Möbius et al. 2014; Henckel et al. 2018; Coutts 2019). Using additional variation in the online experiment, we make some progress toward a characterization of the form of nuisance neglect across information structures. Our results indicate that rather than the (possibly implicit) use of a distorted distribution of $Y$ or a non-Bayesian updating rule, nuisance neglect is best characterized as a strong form of ignorance about the existence of $Y$: People seem to apply a modified signal structure $S_i$ that excludes $Y$.

*3.2.2. External Validity: Nuisance Neglect in a Naturalistic Setting.* While the baseline experimental paradigm provides evidence in a tightly controlled setup, it lacks the ecological validity of real-life contexts in which people typically encounter inference problems. To address this issue and examine the generalizability of nuisance neglect in more naturalistic settings, we designed additional, pre-registered experiments. This series of experiments (i) relies on more real-world contexts that subjects may have some familiarity with, (ii) is not limited to an abstract situation that "feels like a math problem", (iii) leverages one application with more immediate economic relevance, and (iv) includes a version where subjects take a simple choice rather than state a belief given a complex scoring rule.

*Design.* We preserve the basic structure of the baseline experiments, with two random variables ($X$ and $Y$) that causally affect a third variable ($S$), but simplify by binarizing all three variables. We specify the base rates of as well as the causal relationships between the two generative causes ($X, Y$) and the effect variable ($S$). Each subject took decisions in two naturalistic contexts (see complete instructions and decision screens in Online Appendix G.6). In context *Earnings*, a hypothetical company makes a quarterly earnings announcement, which either surpasses or falls short of an analyst prediction. In this scenario, there are two generative causes: the company's operational performance and the general business climate. When the company exceeds the operational performance goal, this *causes* realized earnings to surpass analyst expectations with probability 70%, irrespective of the business climate (see below for a discussion of how to model posterior beliefs given such a causal structure). Conversely, when the business climate is good, this *causes* realized earnings to surpass analyst expectations with probability 90%, irrespective of operational performance. Exceeding operational performance and good business climate are independent of one another and each occur with 50% probability. Subjects then find out that the company's earnings actually surpassed the analyst prediction. Given all this information, a Bayesian would infer that there is a 65% chance that the company exceeded the operational performance goal and a 73% chance that the business climate was good. Similar to the baseline experiment, our main interest was in the treatment comparison between *Broad*, in which subjects stated their beliefs about both operational performance and business climate, and *Narrow*, in which subjects were only asked about operational performance.

The second vignette, *Restaurant*, leverages a context that plausibly taps into subjects' real-life experience. Subjects were asked to imagine having dinner at a new restaurant. The dining experience either exceeds or fails their expectations based on similar restaurants. The actual restaurant quality (which is outstanding or not with equal probability) causes the dining experience to exceed expectations with probability 95%, and good luck that is unrelated to restaurant quality—such as a good mood, sunny weather, or enjoyable company (which happens with 50% probability)—causes an exceeding dining experience with probability 80%. The corresponding Bayesian posterior was 71% for outstanding restaurant quality and 65% for good luck.

*Treatment Conditions and Outcomes.* We ran a total of six between-subjects treatment conditions that cross the main treatment manipulation (*Broad* vs. *Narrow*) with different outcome measures: In *Belief Probabilistic*, subjects received all probabilistic information that was necessary to form a Bayesian posterior similar to the baseline experiment, and were incentivized using a binarized scoring rule. In *Action*, subjects took an action instead of stating a belief. For one (in *Narrow*) or both (in *Broad*) causes, they were endowed with $1 each and could either keep this money or bet it on the occurrence of the respective cause. If the cause occurred, then their $1 bet would be tripled, and if the cause did not occur, then the money would be lost. In *Belief Simple*, we replaced all numerical probability information with verbal descriptions (e.g. 95% was described as an "extremely high" probability and 80% as a "high" probability) and only asked subjects to indicate which state of the target cause (in *Narrow*) or of both causes (in *Broad*) they thought was more likely, without any monetary incentive to avoid the added complexity of a scoring rule. See all the decision screens in Online Appendix G.6.

The rationale behind this series of treatments is that while the first outcome variant (*Belief Probabilistic*) is closest in spirit to a simplified version of the baseline experiments and mainly adds a naturalistic cover story, the second variant (*Action*) removes the artificiality of stating a belief (and the corresponding complex scoring rule), and the third variant (*Belief Simple*) simplifies even further by removing all explicit probabilistic information and the incentive scheme.

*Procedures and Pre-registered Predictions.* The vignettes specified the likelihood with which each of the causes changes the state of the effect; the so-called "causal power" (see e.g. Cheng 1997). The standard Boolean Noisy-Or parameterization for non-deterministic disjunctive interactions between causes of an effect allows us to specify the normative equations for causal inference (for details, see e.g. Pearl 2014). We pre-registered two types of predictions.[20] First, we predicted a treatment effect of the main manipulation, that is, that subjects are more likely to believe that the target cause occurred (or bet on its occurrence) in *Narrow* than in *Broad*. Second, we predicted that the point belief in condition *Narrow* for subjects facing the *Belief*

20. See https://aspredicted.org/w2qi8.pdf.

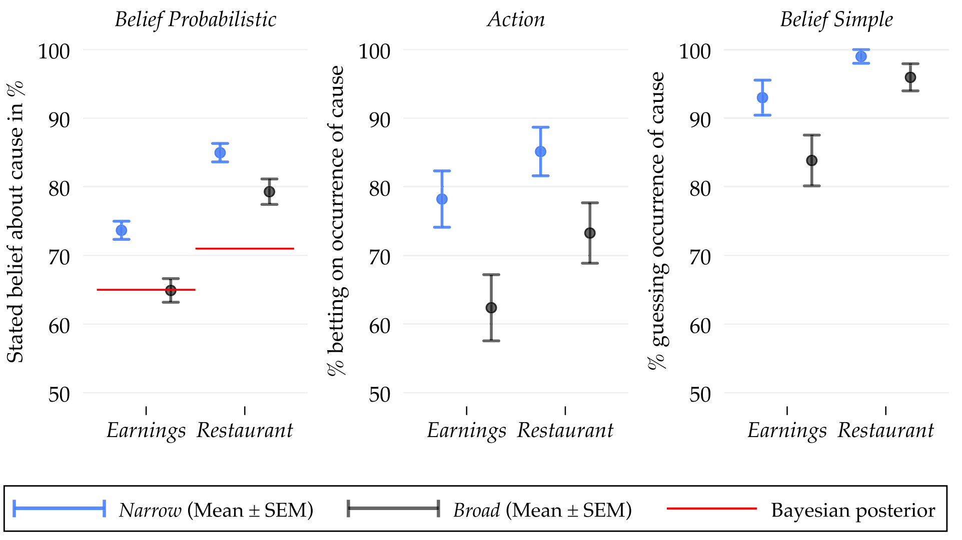


FIGURE 3. Results from experiments with naturalistic task framing. Participants were randomly assigned to either state a probabilistic belief (*Belief Probabilistic*, $N = 199$), take an action by placing a bet (*Action*, $N = 202$), or guess the realization of the focal variable (*Belief Simple*, $N = 199$). Within each outcome group, subjects were randomly assigned to either condition *Narrow* or *Broad*. Each participant completed both the *Earnings* and the *Restaurant* vignettes in random order. Displayed are the mean decisions and standard errors of the mean. The sample size, the treatment effect (*Narrow* vs. *Broad*) and the deviation of the point belief in *Narrow* from the Bayesian posterior in *Belief Probabilistic* were pre-registered.

*Probabilistic* condition would be significantly higher than the Bayesian posterior (as implied by a Noisy-Or model).[21] We pre-registered a total sample of 600 completed responses across the six treatment conditions. A (pseudo-) representative online sample of the US population was collected on Prolific in September 2021.

*Results.* Figure 3 illustrates the results from all six experimental conditions, separately for each of the vignettes. We confirm both pre-registered predictions. First, we find a treatment effect between *Narrow* and *Broad* across all six vignette-outcome pairs.[22] Second, we predicted that beliefs in *Narrow* significantly exceed the Bayesian benchmark for *Belief Probabilistic*, which is also confirmed. Finally, we observe that *Broad* beliefs in *Belief Probabilistic* are indistinguishable from the Bayesian posterior in *Earnings*, but not in *Restaurant*. We did not pre-register a prediction about point beliefs in *Broad*. The reason is that there are a multitude of potential explanations for non-Bayesian point beliefs that vary across the vignettes. For example, point beliefs are likely affected by the idiosyncratic features of the real-life applications. In sum, this

21. Note that there are no Bayesian point predictions for the two other types of outcomes.

22. *Belief Probabilistic*: one-sided $t$-tests yield $p < 0.001$ in *Earnings* and $p = 0.007$ in *Restaurant*. *Action*: two-sample tests of proportion yield $p = 0.007$ in *Earnings* and $p = 0.019$ in *Restaurant*. *Action*: two-sample tests of proportion yield $p = 0.022$ in *Earnings* and $p = 0.085$ in *Restaurant*.

series of treatments strongly supports the external validity of the baseline experiments in more naturalistic task settings.

## 4. Cognitive Mechanisms: Mental Gap or Friction?

### *4.1. Conceptual Considerations*

Research in behavioral economics has produced a collection of deviations from rationality in information processing, many of which are studied and modeled in isolation. Understanding the mechanisms behind updating errors can help identify any common primitives of different anomalies, potentially advancing the convergence of models (Fudenberg 2006) and informing the design of interventions that target specific mechanisms with what Handel and Schwartzstein (2018) label "mechanism policies".

Previous research classifies the sources of deviations from optimality into different categories. We adopt the distinction between "mental gaps" and "frictions" (Handel and Schwartzstein 2018) here as a productive organizing structure that is representative of other, similar classifications. First, a friction occurs if people understand a problem correctly but do not accurately execute all necessary steps to arrive at the normatively optimal solution due to, for example, computational errors, noisy processing, or limited attentional capacity. The corresponding class of models, which includes rational inattention frameworks, assume that belief formation reflects cost–benefit considerations in the presence of some psychological cost of processing information.

Second, a mental gap occurs if people develop an incorrect understanding of the situation to begin with, so that non-Bayesian reasoning is due to *how* they approach and think about a problem given costs. A recent strand of literature in economics examines the implications of misspecified mental models.[23] This work builds on a prominent theme in cognitive science that studies people's mental representations, that is, their subjective models of a problem (Newell and Simon 1972; Fodor and Pylyshyn 1988; Clark 2013; Pitt 2018). Misspecified subjective representations have been characterized through their automaticity, that is, they emerge quickly and effortlessly, they tend to be simple, low-complexity models, and it requires some form of cue to trigger a different representation. This notion of default mental models is related to the intuition-based "System 1" that provides automatic, effortless responses to problems according to dual-process theories (Kahneman 2003; Evans and Stanovich 2013). Dual-process theories also feature the idea that System 1 overrides by the deliberate, effortful System 2 does not occur automatically but requires situational cues (Kahneman 2003; Stanovich and West 2008).

23. See, for example, Schwartzstein (2014), Esponda and Pouzo (2016), Eyster and Rabin (2014), Hanna, Mullainathan, and Schwartzstein (2014), Enke and Zimmermann (2019), Bohren (2016), Bohren and Hauser (2019), Barron, Huck, and Jehiel (2019), Enke (2020), Gabaix (2014), Heidhues, Kőszegi, and Strack (2018), Spiegler (2016), Gennaioli and Shleifer (2010), and Bushong, Rabin, and Schwartzstein (2021).

This section aims to shed light on whether nuisance neglect is better characterized as rooted in a mental gap or a friction. If a mental gap is at the source of nuisance neglect, then the bias should respond to attentional manipulations that alter how a subject thinks about the task. If nuisance neglect reflects a friction, then its prevalence should depend on the relative size of the benefits and cognitive costs associated with an updating problem.

In the following, we present a sequence of analyses that aim to disentangle a mental gap from a friction explanation. Friction explanations have been widely studied in the literature, which includes psychometric designs used in recent economics literature (e.g. Caplin et al. 2020). Our primary focus is on testing for the presence of a mental gap, as work on this topic has received comparably less attention in previous work (but there are exceptions e.g. Enke 2020). Note that the objective is to test *whether* a mental gap is at the source of the bias in the experimental setup under consideration. It is beyond the scope of this paper to characterize what mental models look like in general.

We proceed in two steps. First, in Section 4.2, we test for a mental gap explanation while trying to hold constant the costs and benefits associated with nuisance neglect. Second, in Section 4.3, we briefly outline our findings from tests of whether nuisance neglect responds to its associated costs while plausibly holding the mental representation that people form constant.

### *4.2. On Mental Gaps*

In the following, we present three tests of a mental gap: an explicit hint at the nuisance variable, an implicit nudge for subjects to reconsider their problem representation, and an analysis of subjects' awareness about their nuisance neglect. Online Appendix Table A gives an overview of all experimental treatments.

*4.2.1. The Effect of a Hint.* Treatment *Hint* only provides incentives for estimating $X$ (similar to condition *Narrow*), but adds a contextual cue that shifts attention to the nuisance variable. On every elicitation screen, subjects see a statement that reads "Also think about the role of $Y$". Note that the hint does not provide direct instructions on how to solve the updating problem. If subjects are aware of the relevance of $Y$ in the updating problem, then this hint should have no effect. Moreover, the hint itself neither changes the cognitive costs associated with accounting for $Y$ in processing the signal, nor the incentive for accuracy.

We conduct online experiments on MTurk following the procedures described in Section D and using the task specifications listed in Online Appendix Table E.1. The results of treatment *Hint* and other mechanism treatments are summarized in Figure 4. For each treatment, we pool data from all five updating tasks and display the fraction of beliefs that can be characterized as Bayesian, nuisance neglect, and signal neglect—a posterior belief that equals the prior, as well the remaining fraction of beliefs that

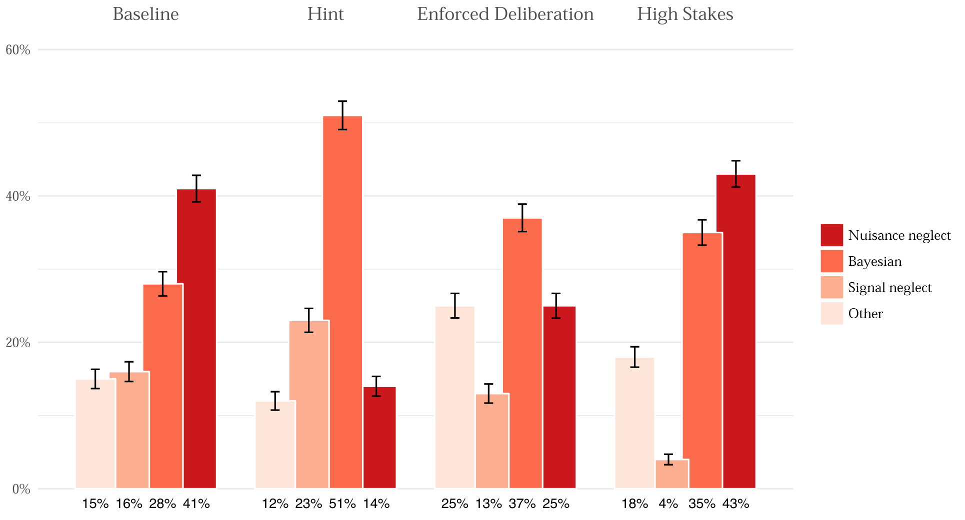


FIGURE 4. Fraction of stated beliefs in line with nuisance neglect, Bayesian updating and signal neglect, as well as the remaining share of beliefs, separately for the baseline online experiment (see tasks specifications in Online Appendix Table E.1) and three mechanisms treatments discussed in Sections 4.2.1 (*Hint*), 4.2.2 (*Enforced Deliberation*), and 4.3 (*High Stakes*). All stated beliefs, pooled across updating tasks. Error bars indicate standard errors of the proportion. Stated beliefs are classified as Bayesian if they are within $\pm 1$ percentage points of the Bayesian posterior, and as nuisance neglect or signal neglect if they exactly corresponded to stating $X = s$ or $X = \mu_X$, respectively.

do not correspond to any of these updating rules.[24] We compare the change in the prevalence of different updating rules relative to the baseline treatment.

Figure 4 documents a substantial and highly significant decrease in the fraction of nuisance neglect by almost two thirds upon adding the hint ($\chi^2$ test, $p < 0.001$). At the same time, Bayesian updating significantly increases from a fraction of below 30% to roughly 50% ($p < 0.001$). Without changing monetary incentives or cognitive costs, the hint has a substantial effect on updating, in support of the idea that subjects are in principle willing and able to update in Bayesian fashion, but fail to notice the need to account for $Y$ to begin with.

*4.2.2. The Effect of Enforced Deliberation Time.* The external hint directs attention to the neglected part of the task. Do subjects notice the nuisance variable on their own if they are nudged to re-consider their solution strategy?

In condition *Enforced Deliberation*, subjects face a 30-second waiting time on each elicitation screen, during which they cannot enter a guess or submit the page. The input

24. Specifically, stated beliefs are classified as Bayesian if they fall into a window of $\pm 1$ percentage points of the Bayesian posterior, and as nuisance neglect or signal neglect if they exactly correspond to stating $X = s$ or $X = \mu_X$, respectively.

fields are only activated after that time is up. This variant of enforced waiting time aims at having people deliberate their approach toward solving the problem—rather than the execution of the subsequent computations—potentially leading them to recognize the need to account for $Y$. Figure 4 shows that this is the case: The share of nuisance neglect in *Enforced Deliberation* falls substantially from 41% to 25% ($p < 0.001$), roughly by half as much as the effect size of a hint.

This result is consistent with an interpretation according to which subjects' solution strategies may be divided into two successive steps: first, parsing the problem description into a mental problem representation, and second, implementing a solution based on that representation. Noticing the neglect may require that subjects specifically reconsider their problem interpretation, rather than their downstream implementation.

*4.2.3. Awareness of Nuisance Neglect.* The effects of a hint and enforced deliberation demonstrate that even minimal interventions that bring attention to $Y$ suffice to substantially reduce nuisance neglect. This indicates that subjects initially fail to think about $Y$ and are thus *unaware* about committing an error. We provide a correlational analysis of this lack-of-awareness hypothesis by measuring confidence in stated beliefs. If subjects who commit nuisance neglect are aware of their distorted beliefs, then they will be less confident than subjects who form optimal beliefs. If instead the simplification of ignoring $Y$ occurs outside of the agent's control, then subjects may deliberately execute the subsequent computations and still exhibit high confidence in their beliefs. In stage *Confidence* that directly follows the belief tasks in the baseline laboratory experiment (Section 3), subjects indicate their willingness-to-accept (WTA) to give up the uncertain payoff associated with each previously stated belief. They are again presented with each individual updating task together with their own stated belief. Participants are asked to indicate whether they prefer to be paid out for the accuracy of their belief or receive a certain monetary amount. They make this binary decision for different fixed amounts, ranging from 0 euros to 6 euros in increments of 0.25 euros. These choices are presented using the multiple-price list method, which Andreoni and Kuhn (2019) argue is particularly easy to understand for subjects. If the task is randomly selected for payout in the end, then a subject's decision in one of the rows of the list is implemented. Note that the *Confidence* tasks (i) have no time limit, such that subjects could freely rethink their stated belief, and (ii) the subjective valuation in each task provides an incentivized measure of confidence in the belief distribution itself, beyond the variance of the stated belief distribution.[25]

Table 3 shows results from regressions in which the dependent variable is the subjective valuation of a stated belief, that is, the minimal certain amount preferred over a having the stated belief paid out. A higher value corresponds to higher confidence in a stated belief. Columns (1)–(3) show that more inattentive beliefs are not significantly associated with lower reservation prices. Even after reconsidering the updating problem

25. The WTA, however, depends on the curvature of the utility function, which motivates robustness analyses below that take into account subjects' risk attitudes.

TABLE 3. Determinants of confidence in stated beliefs.

| Dependent variable: | **Confidence**: valuation for stated belief | | | | |
|---|---|---|---|---|---|
| Condition: | *Narrow* and *Broad* | | | | *Narrow* |
| Estimation method: | OLS | | | IV | OLS |
| | (1) | (2) | (3) | (4) | (5) |
| 0 if *Broad*, 1 if *Narrow* | – 0.497 | – 0.499 | – 0.104 | | |
| | (0.316) | (0.317) | (0.300) | | |
| Inattention $\theta$ | – 0.808 | – 0.801 | – 0.369 | – 0.566 | – 0.487 |
| | (0.509) | (0.508) | (0.512) | (0.429) | (0.436) |
| Treatment dummy $*$ inattention $\theta$ | 0.714 | 0.705 | 0.006 | | |
| | (0.620) | (0.619) | (0.612) | | |
| Variance of belief distribution | | – 0.000*** | – 0.000* | – 0.000* | – 0.000 |
| | | (0.000) | (0.000) | (0.000) | (0.002) |
| Willingness to take risks | | | 0.555*** | 0.557*** | 0.645*** |
| | | | (0.134) | (0.134) | (0.175) |
| Constant | 4.550*** | 4.555*** | 3.694*** | 3.689*** | 3.011*** |
| | (0.180) | (0.181) | (0.637) | (0.619) | (0.613) |
| Task fixed effects | | | Yes | Yes | Yes |
| Additional controls | | | Yes | Yes | Yes |
| $R^2$ | 0.02 | 0.02 | 0.13 | 0.13 | 0.15 |
| # Observations | 1135 | 1135 | 1135 | 1135 | 607 |

Notes: Least-squares and IV regressions. Inattention scores are calculated as $\theta = \frac{H_B}{H_B+H_N}$, where $H_B$ and $H_N$ denote the Hellinger distances of the stated distribution to the Bayesian posterior and the nuisance neglect posterior (as defined in Section 3), respectively. In column (4), implied inattention scores $\theta$ are instrumented with an indicator for treatment status (1 if *Narrow*, 0 if *Broad*). The additional controls include gender, age, income, and task-fixed effect. Robust standard errors clustered at participant level in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

and their own belief, subjects fail to recognize the necessity to account for $Y$ and are equally confident in their own guess. Reassuringly, the variance of the indicated belief distribution is negatively correlated with confidence. While these analyses are correlational in nature, we exploit the causal variation in $\theta$ generated by the treatment manipulation (between conditions *Narrow* and *Broad*) in a regression reported in column (4) of Table 3, which uses treatment status as an instrument for inattention $\theta$. The two-stage least-squares procedure yields a similar coefficient estimate, again indicating no significant relationship between inattention and confidence. Restricting the sample to beliefs stated in *Narrow* (column (5)), we again find no relationship between the valuation of a stated belief and implied inattention.

Taking stock, the effectiveness of simple attentional manipulations and subjects' unawareness of nuisance neglect implied by the confidence measure points to a behavioral mechanism related to how subjects mentally construe the updating problem. Moreover, the findings suggest that selective processing of problem features at least partly depends on factors that are unrelated to the cost of information processing.

Result 2. *Nuisance neglect is reduced by simple attentional cues to $Y$.*

### 4.3. On Cost–Benefit Considerations

Why does the attribution error in inference arise even though it creates a substantial monetary loss? One explanation is that it reflects a simplification strategy that economizes on cognitive resources. A more parsimonious problem representation arguably draws less cognitive capacity, and optimization within a simpler model may allow a quicker and less effortful solution. The avoidance of cognitive effort has been a long-standing theme in cognitive science that has led to the notion of humans as "cognitive misers" or "motivated tacticians" (Stanovich 2009; Fiske and Taylor 2013), with some arguing that most biases in judgment and decision-making reflect effort-reduction strategies (Shah and Oppenheimer 2008). In economics, a growing literature shows that simplifications and inattention can reflect rational, constrained optimization in the presence of cognitive costs or capacity limitations (Sims 2003; Wiederholt 2010; Gabaix 2014; Caplin and Dean 2015). A common prediction of this class of models is that deviations from rationality respond to their associated cost. If nuisance neglect is driven by underlying cost–benefit considerations—explicit or implicit, that is, without the agent's awareness—, then its prevalence should respond to the cognitive costs and the expected benefits of optimal belief updating.

Next, we outline the main findings from examining the effect of variation in the costs of nuisance neglect on its occurrence within the paradigm of this experiment. Put differently, we focus on the sensitivity of updating patterns to changing costs.[26] All details are relegated to Online Appendix E.

In treatment *High Stakes*, the available prize is raised five-fold relative to the baseline online experiment. Under higher incentives, effort as measured by response times increases significantly, both overall and within each subgroup (pairwise $t$-tests, all $p < 0.001$). We find that the prevalence of Bayesian updating increases statistically significantly, but the share of nuisance neglect remains roughly constant. In fact, the increase in Bayesian updating occurs fully at the expense of signal neglect. This means, given higher incentives, subjects try harder, but that only affects non-updating, reducing the fraction of subjects that ignore the signal altogether. On average, higher stakes do not reduce nuisance neglect, however. Compellingly, a ten-fold increase in the stake size in the laboratory experiment leads to a similar pattern (see Online Appendix Section E.4). This indicates that psychic costs, cognitive miserliness, laziness, or effort reduction may explain non-updating, but have limited explanatory power for nuisance neglect.

In Online Appendix E.5, we investigate whether the *specific* monetary cost associated with nuisance neglect affects its prevalence. In economic models of rational belief formation, the likelihood of committing a specific error depends on its expected

26. While the responsiveness of errors to their cost is a central, testable prediction of a friction explanation, a lack of responsiveness does not imply that a mental gap may not in itself be rooted in some form of psychological cost.

cost in utility terms (Wiederholt 2010; Gabaix 2014; Caplin and Dean 2015). On that account, the prevalence of nuisance neglect should vary systematically with its expected loss of accuracy in a given information environment. We vary the expected cost of nuisance neglect using variations of the signal-to-noise ratio and the directional bias implied by the bias. The results suggest that the prevalence nuisance neglect does not systematically respond to these variations in its expected costliness.

## 5. Conclusion

A collection of previous empirical findings implies that misattribution is a pervasive feature of human decision-making. The extant literature, by and large, treats these patterns as unconnected phenomena. This paper contributes in two ways. First, by cleanly documenting the neglect of nuisance variables in both tightly controlled and naturalistic environments, it provides a potential conceptual link between various attribution errors. Because people tend to narrowly focus on explanations that appear subjectively most relevant, they disproportionately assign casual power to these explanations and neglect alternative causes. Second, by studying the precise behavioral mechanisms underlying misattribution, this paper extends the recent literature on updating problems given many pieces of information to attribution problems where agents face a single piece of information. Our conclusion that a similar mechanism of misspecified mental representations may be at play in both problem classes sheds light on the primitives of a theoretical framework that may successfully capture different types of anomalies.

*Limitations and Directions for Future Work.* While this paper provides controlled evidence on the simplest type of attribution problem, it has a number of limitations. First, the combined evidence from over 20 treatments in the abstract and naturalistic paradigms is confined to *static* updating tasks with two random variables. A natural question is how the neglect of alternative causes extends to sequential updating tasks, more complicated signal structures, and environments with more than two variables. Second, the presented evidence from the naturalistic paradigm remains suggestive. This type of extension that transfers structured updating problems to real-world applications merits more work in the future and can help shed light on the relevance and generalizability of belief updating patterns such as nuisance neglect. Third, while this paper concludes that updating errors here may plausibly be due to a mental gap, it does not claim to characterize in any generality when a mental gap occurs, why mental gaps arise, and what they look like in different contexts.

A broader challenge for research on bounded rationality is whether two seemingly conflicting directions in the literature can be reconciled. Evidence on incomplete mental models tends to favor overreaction and "jumping to conclusions", which is broadly in line with the heuristics and biases program and classical work by Kahneman and Tversky. This paper falls into this category. A separate strand of the literature highlights the role of noise and imprecision in human cognition (see Woodford 2019, for an overview). Noisy processing motivates a class of models that predominantly predict insensitivities and underreactions, which is supported by mounting evidence on both lower-level perceptual processes and higher-level reasoning (Steiner and Stewart 2016; Khaw et al. 2019; Enke and Graeber 2022a,b; Frydman and Jin 2022; Gabaix and Laibson 2022). Future work may help shed light on whether these forces operate simultaneously or apply in distinct environments, and if so, what characterizes their respective scope of application.

## References


Abaluck, Jason and Abi Adams (2017). "What Do Consumers Consider Before They Choose? Identification from Asymmetric Demand Responses." NBER Working Paper No. 23566.

Abaluck, Jason and Jonathan Gruber (2011). "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program." *American Economic Review*, 101(4), 1180–1210.

Adam, Klaus, Albert Marcet, and Johannes Beutel (2017). "Stock Price Booms and Expected Capital Gains." *American Economic Review*, 107(8), 2352–2408.

Allcott, Hunt (2011). "Consumers' Perceptions and Misperceptions of Energy Costs." *American Economic Review*, 101(3), 98–104.

Andreoni, James and Michael A. Kuhn (2019). "Is It Safe to Measure Risk Preferences? A Comparison of Four Techniques." Unpublished working paper.

Bandyopadhyay, Prasanta, Gordon Brittan, and Mark Taper (2016). *Belief, Evidence, and Uncertainty: Problems of Epistemic Inference*. Springer.

Bar-Hillel, Maya (1980). "The Base-Rate Fallacy in Probability Judgments." *Acta Psychologica*, 44, 211–233.

Barron, Kai, Steffen Huck, and Philippe Jehiel (2019). "Everyday Econometricians: Selection Neglect and Overoptimism When Learning from Others." WZB Discussion Paper SP II 2019.

Bartoš, Vojtěch, Michal Bauer, Julie Chytilová, and Filip Matějka (2016). "Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition." *American Economic Review*, 106(6), 1437–1475.

Benjamin, Daniel J. (2019). "Errors in Probabilistic Reasoning and Judgmental Biases." In *Handbook of Behavioral Economics*, edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson. Elsevier.

Bertrand, Marianne and Sendhil Mullainathan (2001). "Are CEOs Rewarded for Luck? The Ones without Principals Are." *Quarterly Journal of Economics*, 116, 901–932.

Bohren, Aislinn (2016). "Informational Herding with Model Misspecification." *Journal of Economic Theory*, 163, 222–247.

Bohren, Aislinn and Daniel Hauser (2019). "Social Learning with Model Misspecification: A Framework and a Robustness Result." PIER Working Paper No. 18-017.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer (2020a). "Overreaction in Macroeconomic Expectations." *American Economic Review*, 110(9), 2748–2782.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer (2019). "Diagnostic Expectations and Stock Returns." *The Journal of Finance*, 74, 2839–2874.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018). "Diagnostic Expectations and Credit Cycles." *The Journal of Finance*, 73, 199–227.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2020b). "Memory, Attention, and Choice." *Quarterly Journal of Economics*, 135, 1399–1442.

Brownback, Andy and Michael A. Kuhn (2019). "Understanding Outcome Bias." *Games and Economic Behavior*, 117, 342–360.

Bushong, Benjamin and Tristan Gagnon-Bartsch (2022). "Learning with Misattribution of Reference Dependence." *Journal of Economic Theory*, 203, 105473.

Bushong, Benjamin, Matthew Rabin, and Joshua Schwartzstein (2021). "A Model of Relative Thinking." *Review of Economic Studies*, 88, 162–191.

Caplin, Andrew, Dániel Csaba, John Leahy, and Oded Nov (2020). "Rational Inattention, Competitive Supply, and Psychometrics." *Quarterly Journal of Economics*, 135, 1681–1724.

Caplin, Andrew and Mark Dean (2015). "Revealed Preference, Rational Inattention, and Costly Information Acquisition." *American Economic Review*, 105(7), 2183–2203.

Caplin, Andrew, Mark Dean, and Daniel Martin (2011). "Search and Satisficing." *American Economic Review*, 101(7), 2899–2922.

Caplin, Andrew and Daniel Martin (2015). "A Testable Theory of Imperfect Perception." *The Economic Journal*, 125, 184–202.

Chauvin, Kyle (2020). "A Misattribution Theory of Discrimination." Unpublished working paper.

Chen, Daniel L., Martin Schonger, and Chris Wickens (2016). "oTree–An Open-Source Platform for Laboratory, Online, and Field Experiments." *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Cheng, Patricia W. (1997). "From Covariation to Causation: A Causal Power Theory." *Psychological Review*, 104, 367–405.

Chetty, Raj, Adam Looney, and Kory Kroft (2009). "Salience and Taxation: Theory and Evidence." *American Economic Review*, 99(4), 1145–1177.

Clark, Andy (2013). "Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science." *Behavioral and Brain Sciences*, 36, 181–204.

Coibion, Olivier and Yuriy Gorodnichenko (2012). "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy*, 120, 116–159.

Coibion, Olivier and Yuriy Gorodnichenko (2015). "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review*, 105(8), 2644–2678.

Coutts, Alexander (2019). "Good News and Bad News are Still News: Experimental Evidence on Belief Updating." *Experimental Economics*, 22, 369–395.

De Bondt, Werner F. M. and Richard Thaler (1985). "Does the Stock Market Overreact?" *The Journal of Finance*, 40, 793–805.

Dean, Mark and Nathaniel Neligh (2019). "Experimental Tests of Rational Inattention." Working paper, Columbia University.

Edwards, Ward (1968). "Conservatism in Human Information Processing." In *Formal Representation of Human Judgment*, edited by Benjamin Kleinmuntz. Wiley.

Enke, Benjamin (2020). "What You See is All There is*." *Quarterly Journal of Economics*, 135, 1363–1398.

Enke, Benjamin and Thomas Graeber (2022a). "Cognitive Noise and Intertemporal Choice." NBER Working Paper no. 29577.

Enke, Benjamin and Thomas Graeber (2022b). "Cognitive Uncertainty." NBER Working Paper no. 26518.

Enke, Benjamin and Florian Zimmermann (2019). "Correlation Neglect in Belief Formation." *Review of Economic Studies*, 86, 313–332.

Esponda, Ignacio and Demian Pouzo (2016). "Berk–Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models." *Econometrica*, 84, 1093–1130.

Esponda, Ignacio and Emanuel Vespa (2014). "Hypothetical Thinking and Information Extraction in the Laboratory." *American Economic Journal: Microeconomics*, 6, 180–202.

Esponda, Ignacio and Emanuel Vespa (2019). "Contingent Thinnking and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory." Unpublished working paper.

Evans, Jonathan and Keith E. Stanovich (2013). "Dual-Process Theories of Higher Cognition: Advancing the Debate." *Perspectives on Psychological Science*, 8, 223–241.

Eyster, Erik and Matthew Rabin (2010). "Naive Herding in Rich-Information Settings." *American Economic Journal: Microeconomics*, 2, 221–243.

Eyster, Erik and Matthew Rabin (2014). "Extensive Imitation is Irrational and Harmful." *Quarterly Journal of Economics*, 129, 1861–1898.

Eyster, Erik, Matthew Rabin, and Georg Weizsäcker (2018). "An Experiment on Social Mislearning." Rationality and Competition Discussion Paper Series 73.

Falk, Armin, Anke Becker, Thomas J. Dohmen, David Huffman, and Uwe Sunde (2016). “The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences.” IZA Discussionn Paper No. 9674.
Fernbach, Philip M., Adam Darlow, and Steven A. Sloman (2010). “Neglect of Alternative Causes in Predictive but Not Diagnostic Reasoning.” *Psychological Science*, 21, 329–336.
Fernbach, Philip M. and Bob Rehder (2013). “Cognitive Shortcuts in Causal Inference.” *Argument & Computation*, 4, 64–88.
Fiske, Susan T. and Shelley E. Taylor (2013). *Social Cognition: From Brains to Culture*. Sage.
Fodor, Jerry A. and Zenon W. Pylyshyn (1988). “Connectionism and Cognitive Architecture: A Critical Analysis.” *Cognition*, 28, 3–71.
Frydman, Cary and Lawrence Jin (2022). “Efficient Coding and Risky Choice.” *Quarterly Journal of Economics*, 137, 161–213.
Frydman, Cary and Gideon Nave (2016). “Extrapolative Beliefs in Perceptual and Economic Decisions: Evidence of a Common Mechanism.” *Management Science*, 63, 2340–2352.
Fudenberg, Drew (2006). “Advancing Beyond Advances in Behavioral Economics.” *Journal of Economic Literature*, 44, 694–711.
Gabaix, Xavier (2014). “A Sparsity-Based Model of Bounded Rationality.” *Quarterly Journal of Economics*, 129, 1661–1710.
Gabaix, Xavier (2019). “Behavioral Inattention.” In *Handbook of Behavioral Economics-Foundations and Applications 2*, edited by Douglas Bernheimn, Stefano DellaVigna, and David Laibson. Elsevier, p. 261.
Gabaix, Xavier and David Laibson (2022). “Myopia and Discounting.” NBER Working Paper No. 23254.
Gagnon-Bartsch, Tristan, Matthew Rabin, and Joshua Schwartzstein (2019). “Channeled Attention and Stable Errors.” Working paper, Harvard University.
Gennaioli, Nicola and Andrei Shleifer (2010). “What Comes to Mind.” *Quarterly Journal of Economics*, 125, 1399–1433.
Greenwood, Robin and Andrei Shleifer (2014). “Expectations of Returns and Expected Returns.” *Review of Financial Studies*, 27, 714–746.
Grether, David M. (1980). “Bayes Rule as a Descriptive Model: The Representativeness Heuristic.” *Quarterly Journal of Economics*, 95, 537–557.
Gurdal, Mehmet Y., Joshua B. Miller, and Aldo Rustichini (2013). “Why Blame?” *Journal of Political Economy*, 121, 1205–1247.
Haggag, Kareem, Devin G. Pope, Kinsey B. Bryant-Lees, and Maarten W. Bos (2018). “Attribution Bias in Consumer Choice.” *Review of Economic Studies*, 86, 2136–2183.
Handel, Benjamin and Joshua Schwartzstein (2018). “Frictions or Mental Gaps: What’s Behind the Information We (Don’t) Use and When Do We Care?” *Journal of Economic Perspectives*, 32(1), 155–178.
Hanna, Rema, Sendhil Mullainathan, and Joshua Schwartzstein (2014). “Learning through Noticing: Theory and Evidence from a Field Experiment.” *Quarterly Journal of Economics*, 129, 1311–1353.
Hauser, David J. and Norbert Schwarz (2016). “Attentive Turkers: MTurk Participants Perform Better on Online Attention Checks than Do Subject Pool Participants.” *Behavior Research Methods*, 48, 400–407.
Heidhues, Paul, Botond Kőszegi, and Philipp Strack (2018). “Unrealistic Expectations and Misguided Learning.” *Econometrica*, 86, 1159–1214.
Hellinger, Ernst (1909). “Neue Begründung der Theorie Quadratischer Formen Von Unendlich Vielen Veränderlichen.” *Journal für die reine und angewandte Mathematik*, 136, 210–271.
Henckel, Timo, Gordon Douglas Menzies, Peter G. Moffatt, and Daniel John Zizzo (2018). “Belief Adjustment: A Double Hurdle Model and Experimental Evidence.” CAMA Working Paper No. 1/2018.
Hestermann, Nina and Yves Le Yaouanq (2020). “Experimentation with Self-Serving Attribution Biases.” *American Economic Journal: Microeconomics*, 13, 198–237.

Hossain, Tanjim and Ryo Okui (2013). “The Binarized Scoring Rule.” *Review of Economic Studies*, 80, 984–1001.

Jehiel, Philippe (2005). “Analogy-Based Expectation Equilibrium.” *Journal of Economic Theory*, 123, 81–104.

Kahneman, Daniel (2003). “Maps of Bounded Rationality: Psychology for Behavioral Economics.” *American Economic Review*, 93(5), 1449–1475.

Kahneman, Daniel and Amos Tversky (1982). “On the Study of Statistical Intuitions.” *Cognition*, 11, 123–141.

Khaw, Mel Win, Ziang Li, and Michael Woodford (2019). “Cognitive Imprecision and Small-Stakes Risk Aversion.” NBER Working Paper No. 24978.

Landier, Augustin, Yueran Ma, and David Thesmar (2017). “New Experimental Evidence on Expectations Formation.” CEPR Discussion Paper No. DP12527.

Levy, Gilat and Ronny Razin (2015). “Correlation Neglect, Voting Behavior, and Information Aggregation.” *American Economic Review*, 105(4), 1634–1645.

Malmendier, Ulrike and Young Han Lee (2011). “The Bidder’s Curse.” *American Economic Review*, 101(2), 749–787.

Mankiw, N. Gregory and Ricardo Reis (2002). “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve.” *Quarterly Journal of Economics*, 117, 1295–1328.

Martínez-Marquina, Alejandro, Muriel Niederle, and Emanuel Vespa (2019). “Failures in Contingent Reasoning: The Role of Uncertainty.” *American Economic Review*, 109(10), 3437–3474.

Matejka, Filip and Alisdair McKay (2014). “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model.” *American Economic Review*, 105(1), 272–298.

Matějka, Filip and Alisdair McKay (2015). “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model.” *American Economic Review*, 105(1), 272–298.

Möbius, Markus M., Muriel Niederle, Paul Niehaus, and Tanya S. Rosenblat (2014). “Managing Self-Confidence.” NBER Working paper.

Moore, Don A., Elizabeth R. Tenney, and Uriel Haran (2015). “Overprecision in Judgment.” *The Wiley Blackwell Handbook of Judgment and Decision Making*, edited by Gideon Keren and George Wu. Wiley, pp. 182–209.

Morewedge, Carey K. and Daniel Kahneman (2010). “Associative Processes in Intuitive Judgment.” *Trends in Cognitive Sciences*, 14, 435–440.

Newell, Allen and Herbert A. Simon (1972). *Human Problem Solving*. Prentice-Hall.

Pearl, Judea (2014). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Elsevier.

Pitt, David (2018). “Mental Representation.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, winter 2018 ed. Metaphysics Research Lab, Stanford University.

Rabin, Matthew (2013). “Incorporating Limited Rationality into Economics.” *Journal of Economic Literature*, 51, 528–543.

Rabin, Matthew and Joel L. Schrag (1999). “First Impressions Matter: A Model of Confirmatory Bias.” *Quarterly Journal of Economics*, 114, 37–82.

Ross, Lee (1977). “The Intuitive Psychologist and His Shortcomings: Distortions in the Attribution Process.” In *Advances in Experimental Social Psychology*, Vol. 10. Elsevier, pp. 173–220.

Schwartzstein, Joshua (2014). “Selective Attention and Learning.” *Journal of the European Economic Association*, 12, 1423–1452.

Shah, Anuj K. and Daniel M. Oppenheimer (2008). “Heuristics Made Easy: An Effort-Reduction Framework.” *Psychological Bulletin*, 134, 207–222.

Shiller, Robert (1981). “Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?” *American Economic Review*, 71(3), 421–436.

Sims, Christopher (2003). “Implications of Rational Inattention.” *Journal of Monetary Economics*, 50, 665–690.

Sims, Christopher (2006). “Rational Inattention: Beyond the Linear-Quadratic Case.” *American Economic Review*, 96(2), 158–163.

Sloman, Steven A. and David Lagnado (2015). “Causality in Thought.” *Annual Review of Psychology*, 66, 223–247.

Spiegler, Ran (2016). “Bayesian Networks and Boundedly Rational Expectations.” *Quarterly Journal of Economics*, 131, 1243–1290.
Spiegler, Ran (2017). ““Data Monkeys”: A Procedural Model of Extrapolation from Partial Statistics.” *Review of Economic Studies*, 84, 1818–1841.
Stanovich, Keith E. (2009). *What Intelligence Tests Miss: The Psychology of Rational Thought*. Yale University Press.
Stanovich, Keith E. and Richard F. West (2008). “On the Relative Independence of Thinking Biases and Cognitive Ability.” *Journal of Personality and Social Psychology*, 94, 672–695.
Steiner, Jakub and Colin Stewart (2016). “Perceiving Prospects Properly.” *American Economic Review*, 106(7), 1601–1631.
Taubinsky, Dmitry and Alex Rees-Jones (2018). “Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment.” *Review of Economic Studies*, 85, 2462–2496.
Tversky, Amos and Daniel Kahneman (1971). “Belief in the Law of Small Numbers.” *Psychological Bulletin*, 76, 105–110.
Tversky, Amos and Daniel Kahneman (1983). “Extensional versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment.” *Psychological Review*, 90, 293–315.
Wiederholt, Mirko (2010). “Rational Inattention.” In *The New Palgrave Dictionary of Economics*, edited by Matias Vernengo, Esteban Perez Caldentey, and Barkley J. Rosser, Jr., Springer.
Wolfers, Justin (2002). “Are Voters Rational?: Evidence from Gubernatorial Elections.” Research Paper No. 1730, Graduate School of Business, Stanford University.
Woodford, Michael (2019). “Modeling Imprecision in Perception, Valuation and Choice.” NBER Working Paper No. 26258.

## Supplementary Data

Supplementary data are available at *JEEA* online.