

Cognitive Economic Curves: Quantifying the Clarity of Policy and Consumer Communications

Stefan Bucher*

University of Cambridge

Andrew Caplin

New York University

Katherine Chang

Federal Trade Commission

Ryan Oprea

University of California, Berkeley

Eric Spurlino

Federal Trade Commission

August 2025

Abstract

Regulators and firms alike are increasingly required to ensure that consumer-facing communications—including, but not limited to, policies and disclosures—are *clear and understandable*. Yet there is no rigorous, behaviorally grounded metric to quantify clarity. We introduce *cognitive economic curves*, an empirically estimable framework that ranks any form of communication by its clarity, as revealed through its impact on decision quality and welfare. These curves plot the probability of optimal decision-making as monetary stakes increase, yielding a quantitative measure of the cognitive costs imposed by an information presentation. In incentivized experiments with privacy disclosures—a canonical case of complex jargon—we demonstrate how succinct, focused language significantly improves decision quality. Moreover, distinct forms of complexity have distinct effects: extraneous details lead to the rejection of beneficial options, while verbosity increases acceptance of poor ones. The framework generalizes to disclosures and policies of any kind, financial product terms, and platform or marketing communications, offering a scalable, regulator-ready method for

*Corresponding author. E-Mail: sfb41@cam.ac.uk

We are grateful to participants at the 2nd FTC Conference on Marketing and Public Policy, and the 2nd Workshop on Computational Complexity and Economic Decision Making at the University of Cambridge for helpful comments and suggestions. We thank Yi Xuan (Leo) Zhu for invaluable research assistance. The views expressed in this paper are the authors' own and not necessarily those of the Federal Trade Commission or of any individual Commissioner.

compliance testing and a quantitative optimization tool for firms and policymakers. By linking communication design directly to welfare outcomes, cognitive economic curves provide a quantitative foundation for improving how information is presented to consumers.

Keywords: Cognitive Economics, Disclosure, Regulatory Complexity.

JEL Classification: C91, D18, D83, D91, K40.

1 Introduction

For policies to be effective, they must be *comprehensible*. Laws must be understandable to be obeyed; taxes and subsidies must be transparent to influence behavior; disclosures must be clear to protect consumers; and social programs must be intelligible to be accessed. The same is true for firm–consumer communication: product descriptions, financial offers, platform terms, and marketing claims all rely on the recipient understanding them well enough to act. In both public policy and private markets, clarity is a first-order determinant of behavior.

Yet real-world policy language is often anything but clear. Regulations are frequently dense, legalistic, and difficult to parse, even for highly educated audiences. Disclosures are often choked with extraneous detail, embedded in complex logic, and written in inaccessible prose. In the private sector, firms face parallel challenges: complex pricing terms, loyalty program rules, or subscription conditions may meet formal disclosure requirements while leaving consumers confused or disengaged.

Governments have long recognized the importance of clear communication. In the United States, the 2010 *Plain Writing Act* mandated that federal agencies write documents that are “clear, concise, well-organized, and follow other best practices appropriate to the subject or field and intended audience.” A 2019 White House Memorandum on “Improving Clarity and Transparency in Federal Guidance” urged agencies to improve accessibility and readability. At the state level, some jurisdictions have implemented formal readability requirements. For example, New York Insurance Law (2014) §3102 and Florida Statutes (2024) §627.4145 require insurance policies to achieve a minimum score on the Flesch Reading Ease Test (Flesch, 1948), a simple but widespread metric based only on sentence and word length that is also built into popular software like Microsoft Word and Grammarly.¹

¹See <https://support.microsoft.com/en-gb/office/get-your-document-s-readability-and-level-statistics-85b4969e-e80a-4777-8dd3-f7fc3c8b3fd2> and <https://www.grammarly.com/blog/product/readability-scores/>.

Yet these readability scores focus narrowly on text features and are crude proxies for what regulators actually care about: whether recipients *understand* information and can make informed decisions on that basis. They are insensitive to meaning, context, and decision relevance, and are detached from the actual behavioral or welfare consequences of misunderstanding. The result is that simplification efforts—whether in government agencies or private firms—are often ad hoc, guided by intuition and folk wisdom rather than rigorous, behavioral measures of understanding.

The consequences are not abstract. Confusing health plan documents lead consumers to forgo coverage or make costly choices (Bhargava et al., 2017). Misleading mortgage disclosures contributed to the 2008 financial crisis. Overloaded privacy policies fuel disengagement, numbing the public to real risks (Solove, 2021). Even simple tax incentives go unused because of administrative opacity. Perhaps no policy domain better illustrates the gap between formal disclosure compliance and actual understanding than Medicare Part D. Despite well-intentioned efforts to present options clearly, beneficiaries consistently struggle to choose cost-minimizing plans. A large literature has shown that many enrollees leave hundreds of dollars on the table due to confusion about premiums, formularies, and copayments (Abaluck and Gruber, 2011; Kling et al., 2012). These are not isolated mistakes—they reflect systematic limitations in the comprehensibility of the information environment itself.

In the United Kingdom, the Financial Conduct Authority’s new *Consumer Duty* provides a more meaningful benchmark, shifting the focus from whether information is disclosed to explicitly establishing *consumer understanding* as the desired outcome. The duty requires financial services “firms’ communications to support and enable consumers to make informed decisions about financial products and services [...and...] to be given the information they need, at the right time, and presented in a way they can understand” (Financial Conduct Authority, 2022c, ch. 8). The rules explicitly require that “a firm must test communications before communicating them to retail customers” (Financial Conduct Authority, 2022a, FCA Handbook PRIN 2A.5.10 (1a)), with guidance specifying that “[t]his testing should check communications can be understood by customers, so they can make effective decisions and act in their interests” (Financial Conduct Authority, 2022b, 8.40).² Where firms test communications to maximize sales, the FCA expects “testing capabilities of an equivalent standard” for consumer understanding (Financial Conduct Authority, 2022c, 8.11). “As part of testing, firms should consider which questions or

²The policy also acknowledges a tension with the complexity of existing disclosure requirements: “Where firms must communicate complex information to comply with other disclosure requirements, they should consider what additional steps they can take to support consumer understanding” (Financial Conduct Authority, 2022c, 8.7).

outcomes will elicit accurate measures of understanding in an objective way” (Financial Conduct Authority, 2022b, 8.49).

Meeting such obligations requires a framework that is predictive, quantitative, and behaviorally valid—a tool that can be used both to comply with regulatory requirements and to optimize how information is presented to consumers.

This paper develops such a framework. We introduce *cognitive economic curves*: empirical mappings from incentive strength to decision quality that reveal how easily individuals internalize and act on information. These curves plot decision quality as a function of monetary stakes, with steeper slopes indicating greater clarity and flatter slopes revealing confusion or complexity. Building on Caplin et al. (2020), we show how these curves can be used to generate revealed-welfare measures of comprehensibility. The framework is directly estimable from controlled experiments or large-scale A/B tests, making it suitable for both academic research and operational deployment in firms. Our framework is applicable to any domain where communication clarity matters—public disclosures, tax codes, regulations, financial product terms, platform interfaces, and marketing communications.

To demonstrate its utility, we focus on privacy disclosures: a canonical case where clarity is both legally mandated and widely seen as lacking.³ Privacy policies are long, technical, and often more difficult to read than Kant’s *Critique of Pure Reason* (Litman-Navarro, 2019). It has been estimated that reading all privacy policies one consents to in a year would take between 180 and 300 hours (McDonald and Cranor, 2008).

In an incentivized experiment, we vary two common sources of complexity in such disclosures: (i) *extraneous information*—irrelevant content, and (ii) *verbosity*—long and cumbersome language. We find that succinct, focused disclosures substantially improve decision quality and that these complexity dimensions produce distinct error patterns: extraneous details primarily lead participants to reject beneficial options, while verbosity increases acceptance of poor ones. Cognitive economic curves capture these patterns and quantify the welfare loss from complexity.

By linking communication design directly to measurable behavioral outcomes, our framework delivers on a decades-old goal of regulators such as the US Federal Trade Commission: applying economic tools to assess consumer protection policies. As the FTC observed in the 1980s, “Before this time, decisions dealing with ‘unfair or deceptive acts or practices’ were apparently judged to be issues that would not benefit from an economic perspective. In many ways this is not surprising. The economics of consumer protec-

³“We need to fix GDPR’s biggest failure: broken cookie notices,” Wired. <https://www.wired.co.uk/article/gdpr-cookie-consent-privacy>

tion regulation is essentially contained in the economics of information; twenty years ago there was no ‘economics of information’” (Ippolito, 1986). Our approach finally delivers on that promise, offering regulators and analysts a scientifically grounded tool to assess policy comprehensibility, benchmark competing formulations, and identify high-impact simplifications. More broadly, it connects the economics of information to the practical challenges of designing communications that are both effective and comprehensible.

This approach complements, rather than replaces, other behavioral interventions. While “nudges” have proven effective in shifting choices, behavior change alone does not necessarily improve welfare (see, e.g., Allcott and Kessler, 2019). As Chater and Loewenstein (2023) note, many interventions focus on compliance rather than understanding. Our framework addresses this gap by providing a transparent, incentive-compatible method to assess whether individuals truly understand and can act on information in their own interest.

Section 2 provides some background on the psychometric tradition on which our approach is based, before section 3 lays out our theoretical framework. Section 4 describes our experimental design and section 5 presents our findings. Section 6 discusses the related literature, and section 7 concludes.

2 Cognitive Economic Curves and Psychometrics

Before introducing our novel method of measuring clarity of presentations, it is helpful to sketch the psychometric tradition on which our approach is based. This dates back nearly two hundred years to the work of Ernst Weber (1834) and the introduction of *psychometric curves*. These measure the limits of human perceptual ability and are a fundamental construct of many areas of psychology. A classical psychometric curve, for example, measures the proportion of times that the heavier of two hand-held weights is correctly identified, as a function of the difference in weight.

Such a typical psychometric curve is shown in Figure 1a, plotting a decision-maker’s likelihood of identifying the right weight as heavier against the objective difference between the right and left weights on the horizontal axis. Unlike the experimenter, the participant does not observe the ground truth and instead has to estimate it. The typical sigmoid shape then reflects the empirical observation that the decision-maker sometimes makes mistakes, with a probability that is decreasing in the absolute weight difference: The more similar the two weights are in magnitude, the more difficult it is perceptually not to confuse them.

How does this relate to the challenge of measuring how well decision-makers under-

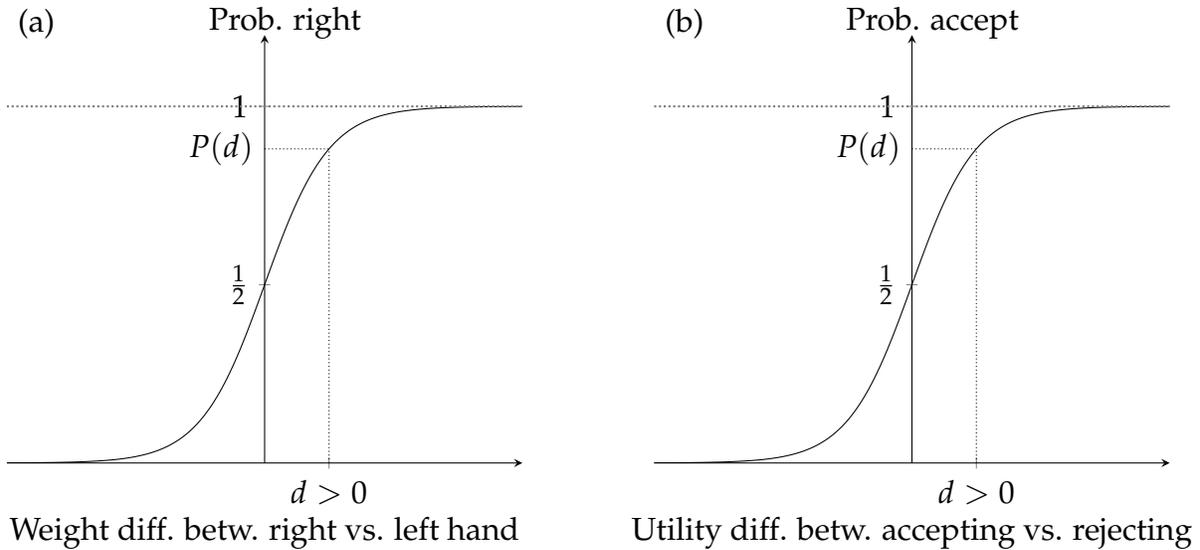


Figure 1: (a) Psychometric curve plotting the probability of identifying the right weight as heavier, as a function of the objective difference between the right and left weights. (b) Cognitive Economic Curve plotting the probability of accepting an option, as a function of the utility difference between accepting and rejecting it.

stand the options they face? The link is indirect but precise. Consider a variation on the perceptual theme in which what has to be identified is not which of two weights is heavier, but rather whether or not a privacy policy is acceptable. To be precise, we consider an experimental design in which the researcher has controlled the set-up so that they in fact know not only whether or not the option at hand is acceptable, but also the utility in which it results.

In such an ideal case, what we might expect conceptually is very much in line with what we see in the case of weight. On the whole, consumers will be more likely to accept the option if it is acceptable, and they will be ever more likely to as the stakes in the decision get larger. So we would expect most to make the correct choice given their values, but also for some to make mistakes. If that is the case, then a curve such as that drawn in Figure 1a might serve to record the proportion of mistakes as a function of the true difference in *value*. Figure 1b draws just such a figure, which we call a *cognitive economic curve*. The horizontal axis measures the utility difference between accepting and rejecting an option, with the vertical axis measuring the probability that the option is accepted. Conceptually the cognitive economic curve is precisely analogous to the psychometric curve except that the underlying *state of the world* represents attributes of the available choice option rather than a standardized percept such as weight.

In abstract terms, one can view both the psychometric experiment based on weights

and the acceptance/rejection decision as having many features in common. First, the answers are *stochastic*: not all answers are identical even given the same objective difference. Second, they depend on the true *state of the world*. In economic applications, the definition of this state of the world is far more subtle and precisely targeted than the weight of an object. It might be the precise terms of a complex contract. It might be the actual contingent cost of insurance as a function of all possible forms of medical problem. It might be the full text of a book. It might be the historical wages of those who receive a certain college degree. In the regulatory settings of this paper, it is what the regulator wants the decision makers to understand. In our primary application it will be the precise nature of the privacy policy. In line with the recent cognitive economic literature, we refer to data of this form as *state-dependent stochastic choice* (SDSC) data.

The big idea is that SDSC data reveal the extent to which decision-makers understand the true state of the world. To give a simple example, when sellers of products engage in obfuscation, their hope is that they will be able to raise prices without customers noticing (Gabaix and Laibson, 2006). Evidence is to be found precisely in the *lack of responsiveness* of demand to changes in this price. Obviously a cognitive economic curve for a binary accept/reject decision is just the first step of many in getting a realistic vision of the value of different modes of presentation, but it is a *sine qua non*. If we cannot make progress in ideal cases, we will be unable to do so in the more intricate settings encountered in practice.

3 Theoretical Framework

This section develops a theoretical foundation for empirically measuring clarity using a cognitive economic approach. Consider a decision-maker (DM) facing a binary choice: accepting an option or rejecting it in favor of an outside alternative. The DM's utility from accepting depends on whether the underlying state of the world $\omega \in \Omega$ belongs to a favorable event E , such as a policy being *good*. It is given by

$$u(\omega; d) = \begin{cases} d & \text{if } \omega \in E \\ -d & \text{if } \omega \notin E \end{cases}$$

with $d \geq 0$ capturing the *stakes* in the decision, while the utility of rejecting is normalized to 0. The DM's prior belief is given by $\mu \equiv \mathbb{P}(\omega \in E)$.

Figure 2 illustrates three different cognitive economic curves for this binary choice problem, plotting the probability $P(u)$ of accepting against the objective utility u of doing

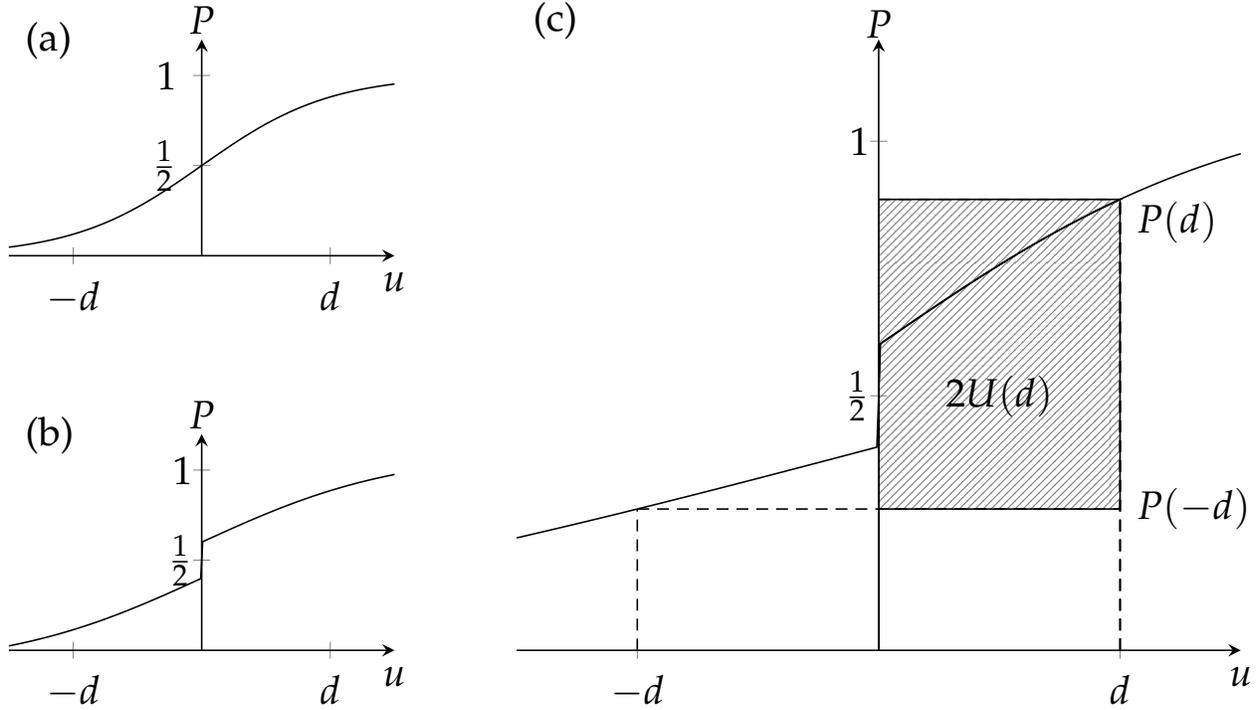


Figure 2: Cognitive economic curves. (a) A cognitive economic curve with a sigmoid shape reminiscent of a psychometric curve. At $u = 0$ the DM is equally likely to take either action. (b) A cognitive economic curve with a discontinuity at $u = 0$, revealing that the DM performs above chance even when the monetary incentive is zero. (c) For this DM, the choice probability is more sensitive to the stakes for good policies than for bad policies. Additionally, for $\mu = \frac{1}{2}$, gross utility $U(d) = \frac{d}{2}[P(d) - P(-d)]$ corresponds to half the area of the shaded rectangle.

so, which is positive or negative depending on the state. Figure 2a shows a curve with a sigmoid shape reminiscent of a psychometric curve, for which the DM is equally likely to accept or reject a policy when the stakes are zero ($d = 0$). In contrast, a DM whose behavior is depicted in Figure 2b is more likely to accept a good policy than a bad policy even when the stakes are zero, with the discontinuity at $u = 0$ revealing that the DM acquires a modicum of information at no cost (possibly because they derive additional non-monetary utility from distinguishing the two).

Figure 2c finally shows a DM whose behavior is more sensitive to the utility of a good policy than the disutility of a bad policy. It further illustrates the area corresponding to the expected gross utility at incentive level d , which is given by

$$U(d) = d [\mu P(d) - (1 - \mu)P(-d)].$$

Note that this is an ex post calculation of realized expected utility resulting from the quality of the choice and as such largely theory-free. If we are willing to make standard assumptions such as rational expectations and that the DM has costs of learning and is rationally inattentive, then we can go much further and use the cognitive economic curve to estimate the costs of learning and therefore also net welfare. The key to doing so is to scale the same challenge by its importance to create a linear class of decision problems in the sense of Caplin et al. (2020). The key construct on which to focus is the *normalized* level of gross utility, which is obtained by dividing expected utility by $d \neq 0$,

$$\bar{U}(d) \equiv \frac{U(d)}{d} = \mu P(d) - (1 - \mu)P(-d) = \underbrace{\mu P(d) + (1 - \mu)(1 - P(-d))}_{A(d)} - (1 - \mu). \quad (1)$$

The normalized utility $\bar{U}(d)$ amounts to the surplus of accuracy $A(d)$ (probability of a correct choice at incentive level d) over the accuracy level $1 - \mu$ attained by rejecting blindly (i.e., opting out of the decision) which yields utility zero.⁴ It is all that is needed to derive costs of learning for rationally inattentive decision-makers using the recovery result of Caplin et al. (2020). The cognitive economic curve has implicit in it the marginal cost of learning, just as a classical supply curve has implicit in it marginal costs of production. A simple proof is in the Appendix; see also Caplin (2023).

Theorem 1. *The cost function for a rationally inattentive DM can be recovered as*

$$K(d) = d\bar{U}(d) - \int_0^d \bar{U}(t)dt.$$

Hence net welfare can be recovered as

$$N(d) = \int_0^d \bar{U}(t)dt.$$

Figure 3 illustrates the welfare recovery result of Theorem 1. The net utility is revealed to be the area between the accuracy curve and $1 - \mu$ vertically and between 0 to d horizontally. The information costs correspond to the area that completes the rectangle extending from $1 - \mu$ to $A(d)$ vertically and from 0 to d horizontally, whose total area reveals the gross utility.

⁴Note that the DM can attain normalized gross utility $\max\{\mu, 1 - \mu\} - (1 - \mu) = \max\{2\mu - 1, 0\}$ at no cost.

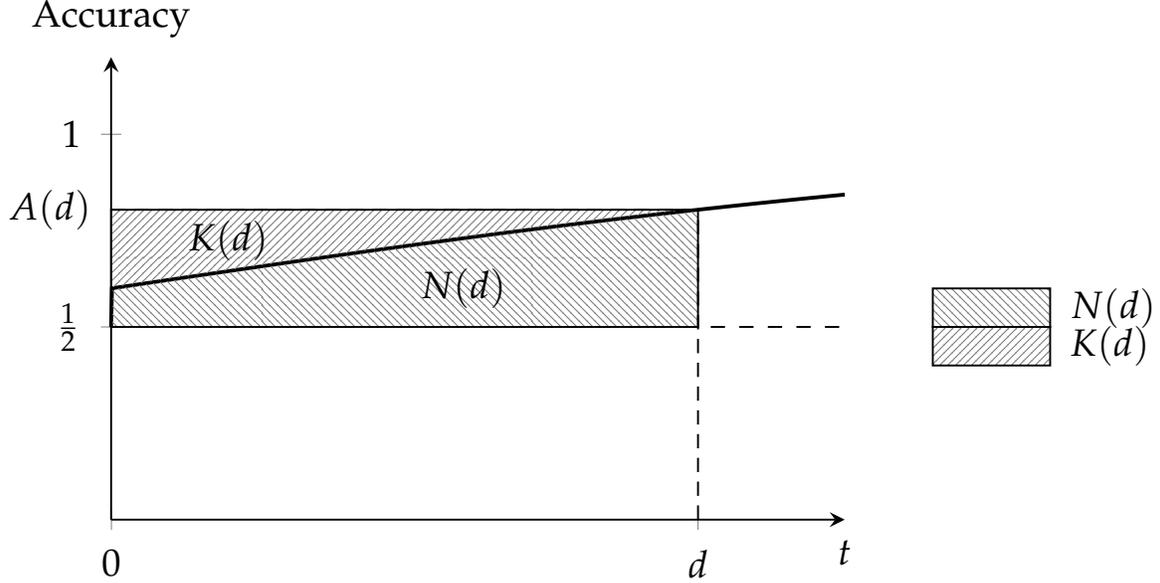


Figure 3: The net utility at incentive level d (for $\mu = \frac{1}{2}$) is revealed to be the area under the accuracy curve $A(t) = \frac{1}{2}P(t) + \frac{1}{2}(1 - P(-t))$ and above $1 - \mu = \frac{1}{2}$, while the information costs are revealed to be the area between the accuracy curve and the accuracy $A(d)$ at d .

The beauty of these figures lies in the fact that they imply a straightforward way to quantify welfare and information costs based solely on choice data at different incentive levels. For a policy maker to assess and compare two candidate disclosures, it thus suffices to conduct an experiment to measure the cognitive economic curves, which then reveal which presentation is clearer: The presentation resulting in a steeper cognitive economic curve will have higher $\bar{U}(d)$ and hence attain higher welfare – it is revealed to be clearer. These revealed welfare measures also allow us to define, in analogy to the aforementioned readability scores, a *comprehensibility score*

$$C := \frac{N(\bar{d})}{\mu\bar{d}} \in [0, 100\%] \quad (2)$$

where $\bar{d} = \max_d d$ is the highest incentive level. This score is on a scale of 0 – 100% because $0 \leq N(d) \leq \mu d$.⁵

While accuracy is a sufficient statistic for welfare in our setting, the cognitive economic curves reveal more than welfare. For example, Figure 4a reveals more about how this accuracy and welfare are attained. At incentive level d , the DM is more likely to falsely accept a bad policy—with probability $P(-d)$ —than they are to falsely reject a

⁵This is due to the fact that $\max\{0, 2\mu - 1\} \leq \bar{U}(t) \leq \mu$ for all $d \geq 0$, which follows from $\max\{\mu, 1 - \mu\} \leq A(d) \leq 1$.

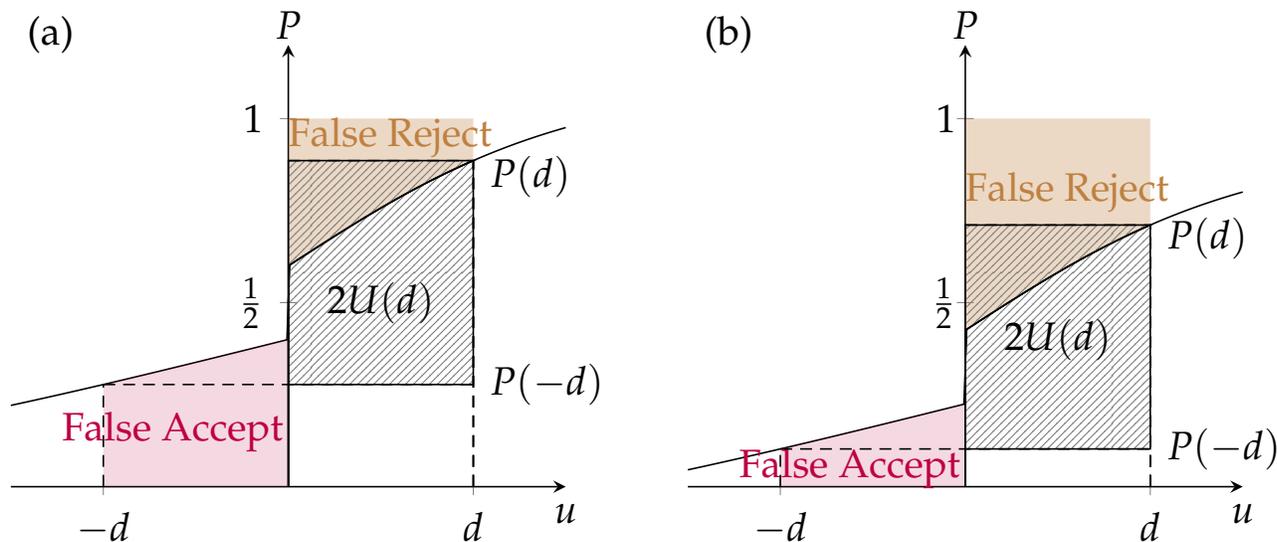


Figure 4: Two cognitive economic curves attaining the same level of accuracy and welfare but with different error types. The curve in (a) makes similarly many false rejections of good policies and false acceptances of bad policies, while the curve in (b) makes frequent false rejections of good policies but few false acceptances of bad policies.

good policy—with probability $1 - P(d)$. The same welfare and level of accuracy could have been attained by a decision-maker exhibiting a different pattern of mistakes altogether (Figure 4b). In Section 5 we will indeed find this to be the case, and discuss in detail how the two complexity dimensions differentially affect the types of mistakes participants make.

These theoretical results provide a precise method to quantify disclosure clarity through direct empirical observation. Cognitive economic curves yield detailed insights into disclosure effectiveness, explicitly distinguishing between cognitive errors. These insights guide targeted improvements in disclosure design and facilitate rigorous, scientifically informed comparisons of clarity.

4 Experiment Design

We illustrate our approach with an experiment whose objective is to empirically elicit cognitive economic curves using real-world decision problems.⁶ We chose the context of privacy policy disclosures for this proof of concept, but our design is easily adapted to numerous other contexts. Our experimental setup closely mirrors the theoretical framework presented previously.

⁶The experiment code is available at <https://github.com/CoCoEcon-Lab/disclosures>.

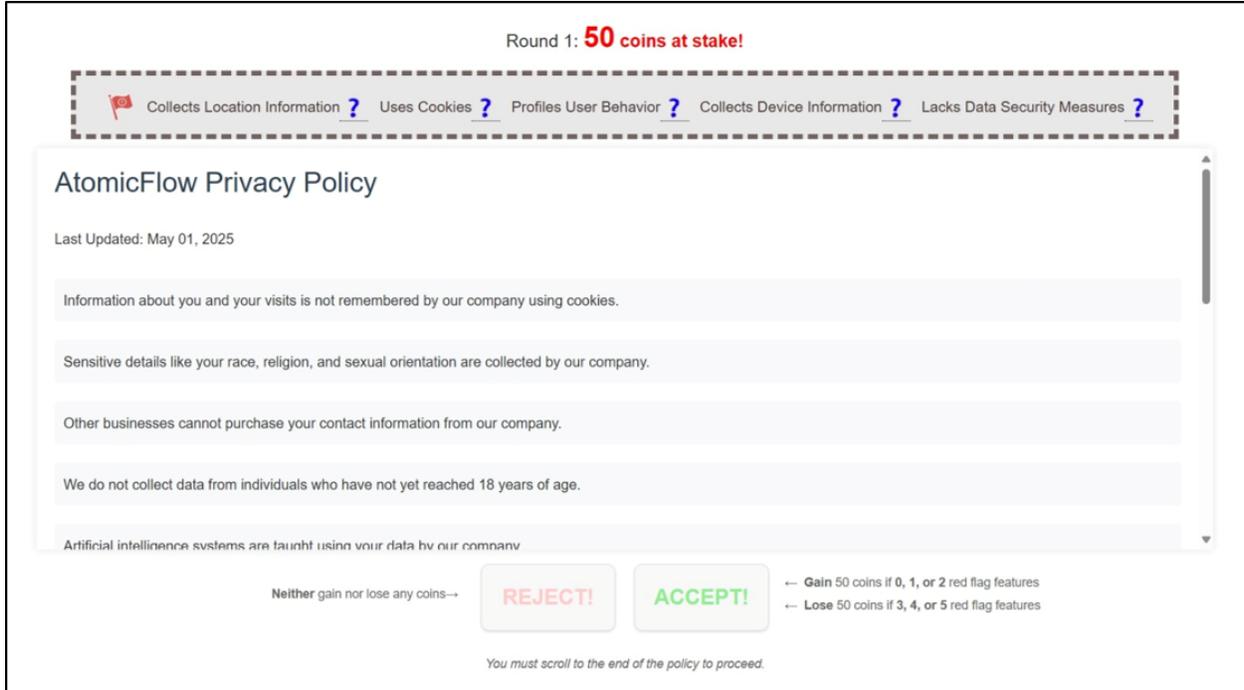


Figure 5: Example choice screen featuring an accept/reject decision on a *succinct* policy with *ten extraneous attributes*. The incentive level d is shown at the top of the screen, and the red flag attributes in the box above the scrollable policy.

Task & Incentives Participants completed 36 trials, each involving a binary accept/reject decision concerning a hypothetical app with a privacy policy $x \in \{0, 1\}^n$ characterized by n binary attributes. Each policy comprised attributes randomly drawn from a grand set of 20 possible attributes (for example, “we [do not] use your device’s GPS or other location services”), and the binary attribute values were drawn from independent Bernoulli distributions with probability 0.5.

On each trial, participants were informed which set of five attributes (among those described in the policy) was to be avoided (“*red flags*”) on that trial. For example, the use of GPS location services could be a red flag to be avoided. Accepting a policy with 2 or fewer red flags present yielded a $(\frac{1}{2} + d)$ -chance of winning a \$8 prize; accepting a policy with more than 2 red flags yielded a $(\frac{1}{2} - d)$ -chance of winning the prize. Rejecting always yielded a 50% chance of winning the prize. The stakes d varied across trials between 0%, 5%, and 50%. The normalized utility of rejecting a policy is thus 0 and the utility of accepting a policy is given by

$$u_{\theta}(x) = \begin{cases} d & \text{if } \theta^T x \leq 2 \\ -d & \text{if } \theta^T x > 2 \end{cases} \quad (3)$$

where $\theta \in \{\{0,1\}^n : \sum_{i=1}^n \theta_i = 5\}$ encodes the trial’s red-flag attributes and \mathbf{x} its policy content.⁷ Note that the experiment’s prior is uniform, since the attributes’ independent Bernoulli distributions imply that $\mu = \mathbb{P}(\theta^T \mathbf{x} \leq 2) = \frac{1}{2}$ for any θ .

Treatments We used a 2×2 within-participant design (see Figure 6) varying (1) the policies’ number of extraneous attributes (0 or 10 irrelevant attributes alongside the 5 red-flag attributes) and (2) whether the attributes were either succinctly described in clearly delineated lines (*succinct*) or embedded within verbose, legalistic paragraphs composed of two to three different attributes (*verbose*).

Participants completed, unbeknownst to them, four trials for each extreme condition (succinct without extraneous items and verbose with extraneous items) and two trials for each intermediate condition at each of the three incentive levels, totaling 36 trials presented in random order.

Instructions and Implementation To implement the aforementioned incentives intuitively, participants were instructed at the beginning of the experiment that they would begin each trial with 50 coins and that they could win or lose $d \in \{0, 5, 50\}$ coins by accepting a policy, or maintain the 50 coins by rejecting it. Each trial’s incentive level d was displayed on a separate screen preceding the choice, and in addition prominently at the top of the choice screen (see Figure 5).

The relevant attributes were designated as “red flags” to be avoided, and displayed prominently above the policy; details on the relevant attributes’ meaning was revealed by hovering the mouse over the associated question marks. Participants were free to spend as much time as they wished before clicking the “ACCEPT!” or “REJECT!” button (the order of which was counterbalanced) below a scrollable text box containing the policy. To ensure that participants had seen the whole policy, they were required to scroll to the bottom of the policy for the response buttons to become active, as is the case for many real-world policies.

⁷Since one of 36 trials was paid out randomly (see below), this utility function on one trial follows from the normalization $u(\$8) - u(\$0) = 36$ so that rejecting a policy yields expected utility

$$\frac{1}{36} \left[\frac{1}{2} u(\$8) + \frac{1}{2} u(\$0) \right] + \frac{35}{36} u(\$0) = \frac{1}{72} u(\$8) + \frac{71}{72} u(\$0) = \frac{1}{2} + u(\$0)$$

and accepting a good policy yields expected utility

$$\frac{1}{36} \left[\left(\frac{1}{2} + d \right) u(\$8) + \left(\frac{1}{2} - d \right) u(\$0) \right] + \frac{35}{36} u(\$0) = \frac{1}{2} + u(\$0) + d$$

and analogously for a bad policy. Normalizing $u(\$0) = -1/2$ yields equation 3.

AtomicFlow Privacy Policy

Last Updated: May 01, 2025

We do not use your behavior and information to build comprehensive user profiles.

Information about you and your visits is not remembered by our company using cookies.

We obtain data about the particular device you use, such as your phone or computer.

We lack robust measures to secure your data against breaches or misuse.

We use your device's GPS or other location services to track your whereabouts.

(a) Succinct $\times 0$ extraneous

AtomicFlow Privacy Policy

Last Updated: May 01, 2025

Information about you and your visits is not remembered by our company using cookies.

Sensitive details like your race, religion, and sexual orientation are collected by our company.

Other businesses cannot purchase your contact information from our company.

We do not collect data from individuals who have not yet reached 18 years of age.

Artificial intelligence systems are taught using your data by our company.

We do not use your behavior and information to build comprehensive user profiles.

We obtain data about the particular device you use, such as your phone or computer.

Personal details such as your age, gender, and income are gathered by our company.

For our European users, our practices do not fully align with European data protection laws (GDPR).

We use your device's GPS or other location services to track your whereabouts.

We lack robust measures to secure your data against breaches or misuse.

We show you targeted ads based on information about your online activities.

We obtain data about you from your social media accounts.

For our California users, our practices do not fully align with California's privacy laws (CCPA).

We disregard the 'Do Not Track' preferences set in your browser.

(c) Succinct $\times 10$ extraneous

AtomicFlow Privacy Policy

Last Updated: May 01, 2025

The Company expressly states that it does not develop detailed user profiles from behavioral patterns or submitted information. This policy of non-profiling encompasses all user activities and data. The Company has established procedures to prevent such profile creation. The Company expressly states that it does not employ cookie technology in any form. This encompasses non-collection of user activity and site visitation data through cookies. The Company has established procedures to prevent cookie implementation.

The Company hereby notifies users that it gathers specific information about their devices, including computers and mobile equipment. This encompasses collection of device identifiers and specifications. The processing of such information follows established procedures. The Company hereby notifies users of the absence of robust information security measures. This encompasses insufficient safeguards against unauthorized access and misuse. The Company operates without comprehensive data protection protocols. The Company hereby notifies users that it systematically collects and processes location data through authorized tracking technologies. This includes information gathered via device GPS and other location services. The collection and retention of such geographic data occurs in accordance with Company procedures.

(b) Verbose $\times 0$ extraneous

AtomicFlow Privacy Policy

Last Updated: May 01, 2025

The Company expressly states that it does not employ cookie technology in any form. This encompasses non-collection of user activity and site visitation data through cookies. The Company has established procedures to prevent cookie implementation. The Company hereby notifies users of its collection and retention of sensitive personal information through authorized channels. This encompasses data including, but is not limited to, information about race, religious affiliation, and sexual orientation. The processing of such sensitive information is governed by internal data handling procedures.

The Company expressly states that it does not sell user contact information to external businesses. This policy of non-sale encompasses all contact detail transfers. The Company has established procedures to prevent such commercial transactions. The Company expressly states that it does not collect information from individuals under eighteen (18) years of age. This policy of non-collection encompasses all minor user data. The Company has established procedures to prevent gathering of such information. The Company hereby notifies users that their data is utilized in artificial intelligence system training. This encompasses the processing of user information for AI development purposes. The implementation of such training follows established protocols.

The Company expressly states that it does not develop detailed user profiles from behavioral patterns or submitted information. This policy of non-profiling encompasses all user activities and data. The Company has established procedures to prevent such profile creation. The Company hereby notifies users that it gathers specific information about their devices, including computers and mobile equipment. This encompasses collection of device identifiers and specifications. The processing of such information follows established procedures. The Company hereby notifies users that it systematically collects and maintains demographic information through authorized channels. Such demographic information includes, but is not limited to, user age, gender, and income-related data. This information is processed and retained in accordance with established Company protocols.

The Company hereby notifies European users of incomplete compliance with GDPR requirements. This encompasses various aspects of data protection regulations. The Company's practices do not fully satisfy GDPR standards. The Company hereby notifies users that it systematically collects and processes location data through authorized tracking technologies. This includes information gathered via device GPS and other location services. The collection and retention of such geographic data occurs in accordance with Company procedures. The Company hereby notifies users of the absence of robust information security measures. This encompasses insufficient safeguards against unauthorized access and misuse. The Company operates without comprehensive data protection protocols.

The Company hereby notifies users that it processes online behavioral data to deliver targeted advertisements. This encompasses the analysis of user online activities for promotional purposes. The implementation of such targeted advertising follows established Company procedures. The Company hereby notifies users of its collection of information from social media platforms and services. This encompasses data gathering from social networking accounts and activities. The processing of such social media information follows established procedures.

The Company hereby notifies California users of incomplete compliance with CCPA requirements. This encompasses various aspects of privacy regulations. The Company's practices do not fully satisfy CCPA standards. The Company hereby notifies users that it disregards browser-based 'Do Not Track' signals. This encompasses continued tracking activities regardless of user preferences. The non-implementation of such signals follows Company procedures.

(d) Verbose $\times 10$ extraneous

Figure 6: Illustration of the four treatments. Succinct disclosures (left) have short sentences in delineated sections for each attribute, while verbose disclosures (right) have long sentences in paragraphs containing two to three attributes each. Additionally, disclosures with zero extraneous attributes (top) have a shorter text screen, while disclosures with ten extraneous attributes (bottom) have a longer text screen. The detailed text of this example is shown in Appendix D.

Participants were informed they would review 12 policies at each of the three incentive levels, and that the coins earned in a randomly selected trial would be converted to raffle tickets (out of a total of 100) determining their probability of receiving an \$8 bonus in a raffle at the end of the experiment, in addition to the base pay of \$5. The median completion time was 50 minutes.

Understanding Prior to the binary choice task, participants were tested on their understanding of the incentives in an understanding test in which participants were asked to report the number of coins earned in twelve scenarios: accept/reject a policy with more/less than two red flags, at each of the three incentive levels. Participants with 10 out of 12 correct answers by their second attempt were allowed to proceed with the experiment (but not before they correctly responded to all 12 questions); 83.3% of those who passed the understanding check answered all questions correctly in their first attempt. After each attempt, participants were notified of which questions were answered incorrectly. In line with Prolific’s policies, participants were able to access the instructions while completing the comprehension test.

Questionnaire At the end of the experiment, participants were asked to complete a short survey collecting basic demographic information along with a cognitive reflection test (Frederick, 2005), as well as their subjective beliefs about their performance under the different treatments and incentive levels.

5 Results

We collected data from $N = 300$ participants (42.6% female; age 39 ± 12) on the online experiment platform Prolific.⁸

Cognitive economic curves The empirical cognitive economic curves from our experiment are presented in Figure 7. The curves systematically differ across treatments, indicating that both extraneous information and verbosity meaningfully impact decision quality. In all treatments, the acceptance probability is non-decreasing in the utility of accepting. All curves exhibit a sigmoid shape reminiscent of classical psychometric curves, albeit with a discontinuous jump at $u = 0$, revealing—as discussed above—that participants distinguished between good and bad policies to an extent even when the monetary incentive was nil.

⁸The data are available at <https://osf.io/vgpbr/>.

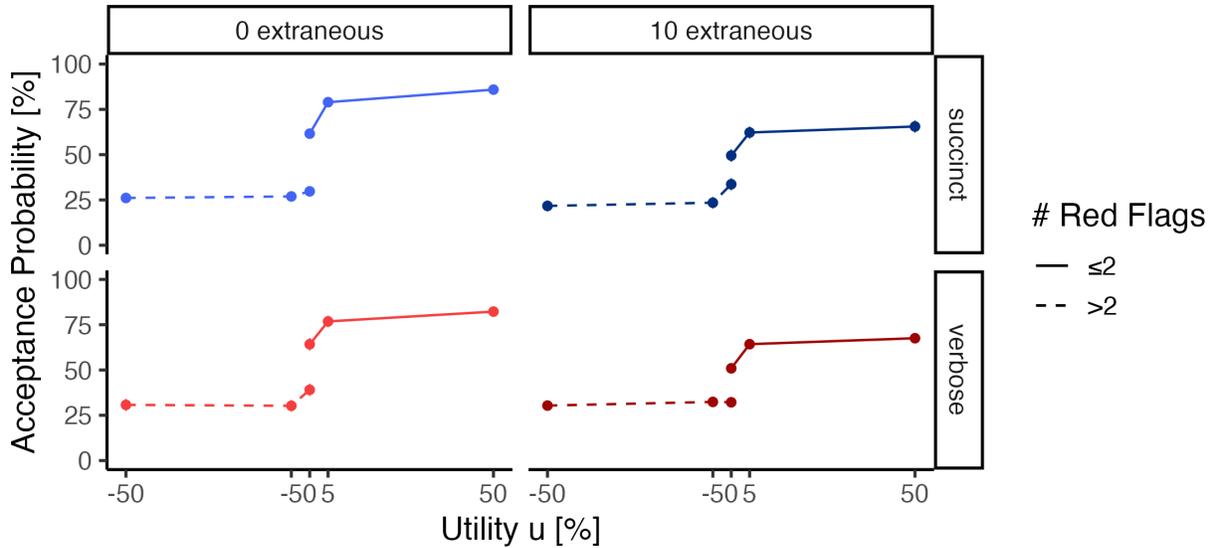


Figure 7: Cognitive economic curves show, for each treatment, the acceptance probability as a function of the utility of accepting relative to rejecting. These are the empirical analogue to Figure 2.

Accuracy The accuracy (probability of a correct choice) implied by these cognitive economic curves is shown in Figure 8. Accuracy is higher when policies contain no extraneous attributes compared to policies with extraneous attributes. Similarly, succinct disclosures outperform verbose disclosures. Notably, succinct disclosures containing extraneous attributes achieved lower accuracy than verbose disclosures without extraneous attributes, indicating that the verbosity of the language has a smaller effect on accuracy than ten extraneous attributes. Accuracy improves with increased stakes across all treatments. The largest gain in accuracy occurs between no stakes and low stakes, while the improvement from low to high stakes is smaller. We find no significant interaction between the treatments and the incentive levels; instead, treatments broadly shift accuracy uniformly across all incentive levels.

Welfare The accuracy curves of Figure 8 (as well as the curves of Figure 16 in the appendix) allow us, according to Theorem 1, to compute the empirically revealed gross welfare $U(d)$ and net welfare $N(d)$ as well as information costs $K(d)$, which are plotted in Figure 9 as a function of the stakes d . Note that welfare is characterized in units of percentages (of $u(\$8)$, the utility of receiving the bonus with certainty); it is bounded above by 25%, which is the *additional* probability of winning \$8, above and beyond the 50% chance associated with rejecting an option, that a decision-maker can afford themselves

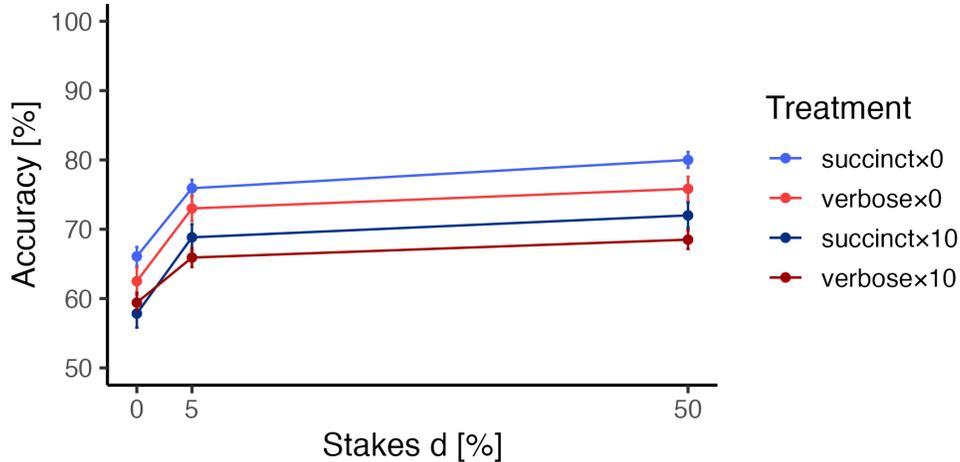


Figure 8: Accuracy across treatments as a function of stakes, analogous to Figure 3.

with perfect choices.⁹

Gross welfare and welfare net of information costs increase with stakes and are highest for succinct disclosures without extraneous information and lowest for verbose disclosures with extraneous information. The information costs incurred on a decision are also increasing in the stakes, and their order of magnitude is one percent of the utility attributed to the \$8 bonus. Note that succinct disclosures with extraneous information incur higher information costs and attain lower net welfare than verbose disclosures with no extraneous information, but surprisingly the lowest information costs are incurred for verbose disclosures with extraneous information, which still yield the lowest net welfare. This demonstrates the rich conclusions that can be drawn from our approach, as average accuracy alone could not have revealed this, and may reflect the fact that participants fail to engage with “walls of text” so that their expenditure on information costs is low despite (or precisely because of) high marginal information costs.

Readability Scores and Individual Differences Our welfare measures can also be used to compute a *comprehensibility score* that has a number of advantages relative to readability scores such as the Flesch (1948) reading ease score. In the context of the experiment, $0 \leq d \leq 50\%$ and $\mu = \frac{1}{2}$ so that the comprehensibility score of equation 2 reduces to

$$C = 4N(50\%). \quad (4)$$

⁹The total winning probability is bounded at 75% because 50% of policies are bad and hence to be rejected which yields a 50% winning chance, and 50% of policies are good and to be accepted which yields a 100% winning chance.

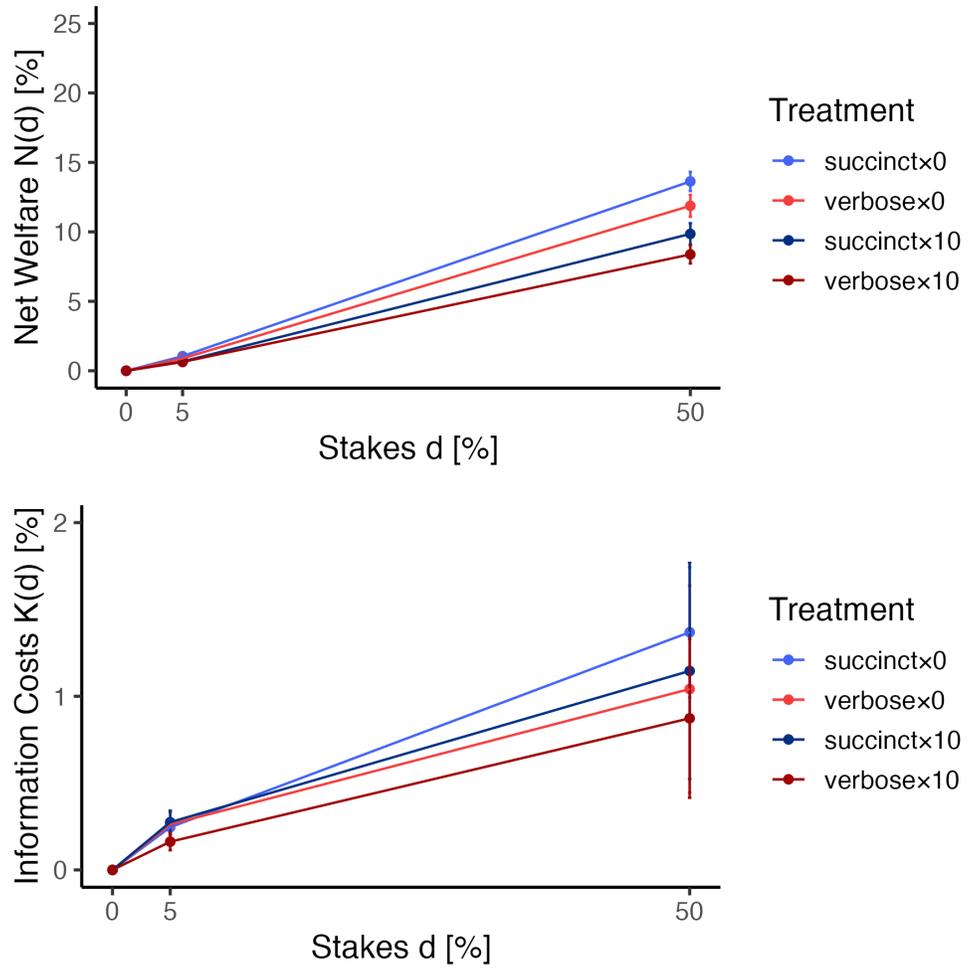


Figure 9: Revealed net welfare and information costs as a function of stakes.

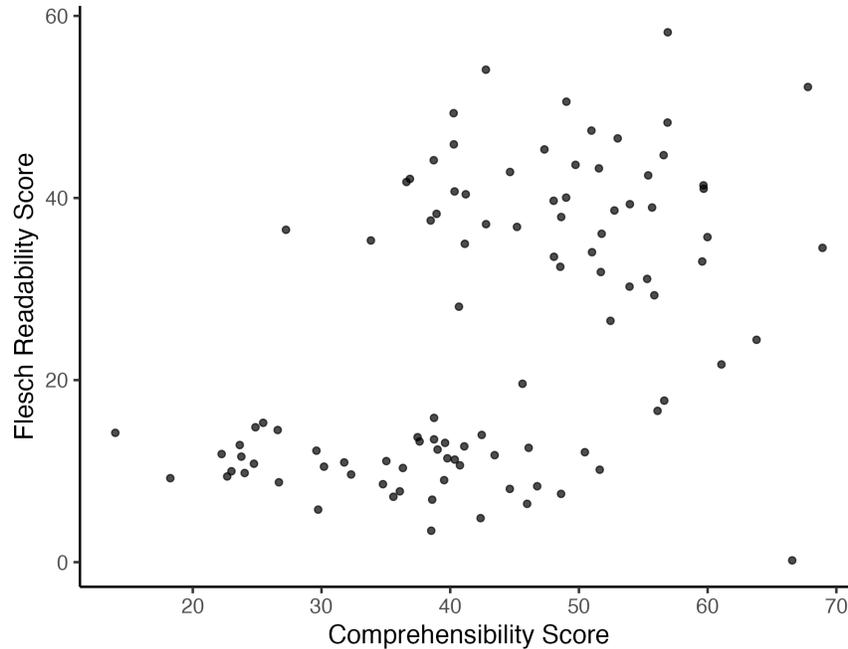


Figure 10: For this scatterplot, the Flesch reading ease scores of all disclosures shown in the experiment were computed and then split into 100 bins containing equal numbers of trials. Since computing welfare measures and our comprehensibility score requires aggregating choices across trials, we computed the comprehensibility score for all policies in each of the 100 bins, and plotted it against each bin’s average Flesch reading ease score.

Although theoretically this score is on a scale of 0 – 100%, empirical readings on the interval $[-100\%, 100\%]$ are possible because on any given sample it is conceivable that a participant reliably makes the inferior choice. A score of 0% means that the choice behavior reveals no signal about the state of the world, while a score of 100% means that the choice behavior perfectly responds to the state of the world.

Figure 10 illustrates the positive correlation between our comprehensibility score and the Flesch reading ease score, supporting the validity of cognitive economic curves as a measure of disclosure comprehensibility. Yet it also makes clear that the scores are far from always being in agreement. Notably, Figure 11 reveals that our welfare-based comprehensibility score is able to rank the four treatments, whereas the Flesch reading ease score responds only to the verbosity of the language but not to the amount of extraneous information, because it does not take the length of a text into account.

Perhaps more importantly, the Flesch reading ease score is purely text-based and hence cannot account for individual differences. These can be substantial, as is illustrated by Figure 12 plotting each participant’s comprehensibility scores across two treatments. The welfare-based comprehensibility scores vary widely across participants, and are cor-

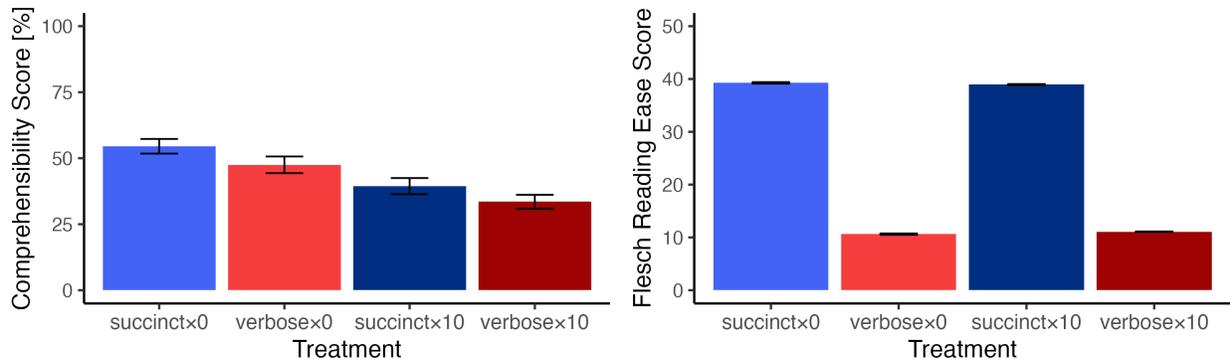


Figure 11: Our welfare-based comprehensibility score and the Flesch reading ease score by treatment.

related across the two treatments shown. The majority of participants (62.3%) are located below the diagonal and hence attain a higher welfare score for succinct disclosures without extraneous information than for verbose disclosures with extraneous information. The fact that for many the opposite is true may in part be due to sampling noise, but it also illustrates the approach’s ability to detect individual differences, which may tell a richer story than population averages and open the door for personalized information policies.

Error patterns Cognitive economic curves reveal more about the cognitive process than the welfare in which it results. For example, they provide insights into the specific types of decision errors participants make. Figure 13 summarizes the distinct impact of extraneous information and verbosity on error rates, displaying the probabilities (aggregated across stakes) with which participants falsely reject a policy with ≤ 2 red flags and with which they falsely accept a policy with > 2 red flags. In the language of statistical hypothesis testing (of the null hypothesis that the policy has ≤ 2 red flags), a false rejection is a type I error (“false positive”) and a false acceptance is a type II error (“false negative”).

Figure 13 reveals a clear pattern of mistakes: Extraneous information increases false rejections but not false acceptances, while verbosity increases false acceptances but not false rejections. This asymmetry suggests that each dimension of complexity affects cognitive processes differently. This nuanced pattern cannot be explained by a shifting acceptance threshold alone, which would affect false rejections and false acceptances simultaneously. Instead, it appears consistent with a simple signal detection theory model in which extraneous information increases the noise with which good attributes are perceived, and verbosity increases the noise with which bad attributes are detected. This account is also consistent with the acceptance probabilities shown in Figure 22 in the appendix, which

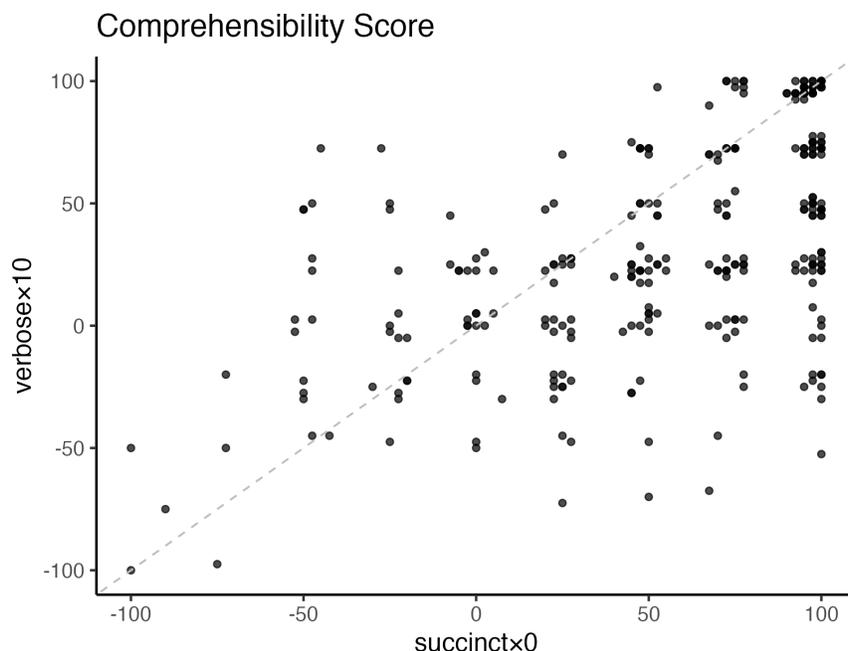


Figure 12: Each dot shows one participant’s comprehensibility scores for succinct disclosures with no extraneous information and verbose disclosures with extraneous information. Participants cover the full range of feasible comprehensibility scores. The scores’ discrete grouping is an artifact of the limited number of trials at each incentive level determining each participant’s score for a given treatment, with behavior at $d = 50$ having a relatively larger impact on welfare than at $d = 5$.

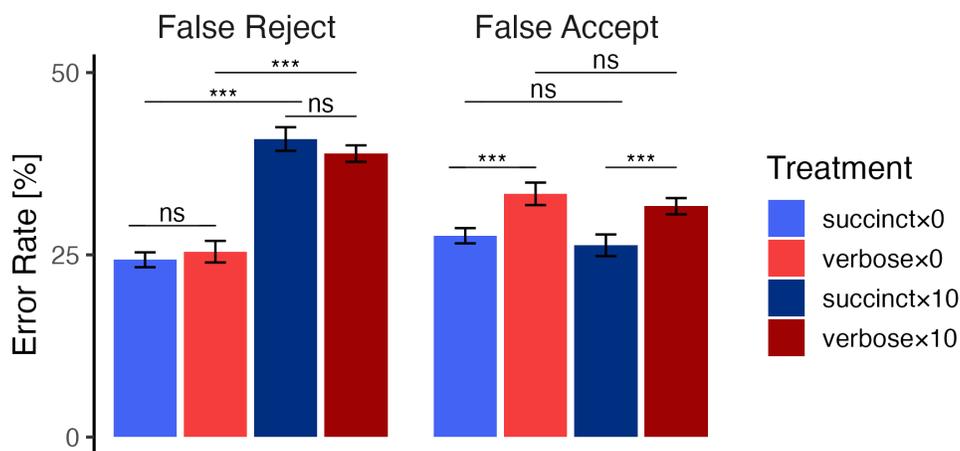


Figure 13: The two treatment dimensions result in different error types: Extraneous information increases false rejections (type I errors; False Positive) but not false acceptances, while verbosity increases false acceptances (type II errors; False Negative) but not false rejections. * * * denotes statistical significance at $p < 0.01$, *ns* indicates difference not significant.

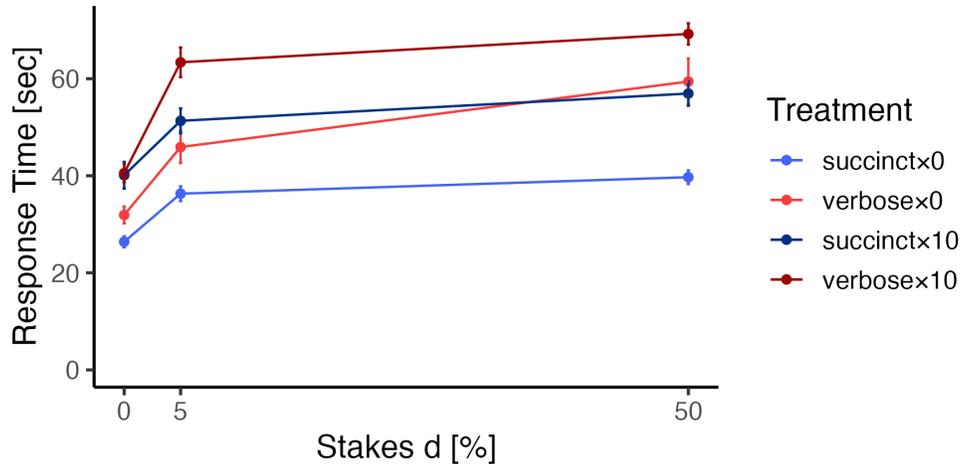


Figure 14: Response times by treatment as a function of stakes.

reveal more rejections in the presence of extraneous information and marginally fewer for verbose compared to succinct policies.¹⁰

The cognitive economic curves further show that false acceptances are much less sensitive to the incentive level than false rejections (see Figure 17), suggesting that the latter can be reduced by raising the stakes while the former cannot.

Response Times Prior work highlighted the real-world time burden imposed by privacy policies (Litman-Navarro, 2019). Figure 14 shows that response times increase with stakes, especially in verbose conditions, suggesting that participants allocate more cognitive effort when incentives are higher. We find that response times are slower with extraneous attributes than without, and slower for verbose treatments than succinct treatments. Notably, average response times are higher for succinct policies with extraneous information than for verbose policies without extraneous information. Yet, despite these longer deliberation times, succinct disclosures with extraneous content result in lower accuracy than verbose disclosures without extraneous content (Figure 3). This pattern suggests that verbosity may be cognitively easier to process than extraneous information, yielding a higher accuracy return per unit of time.

6 Literature Review

Clear and conspicuous disclosure has long been a tool used by regulatory bodies to ensure informed consumer decision-making. Ben-Shahar and Schneider (2014) critically assess

¹⁰We also observed a rejection bias at no stakes, but not for positive stakes (see Figure 23).

mandated disclosures, noting that such disclosures often score extremely poorly on simplistic readability metrics such as grade-level scores. Loewenstein et al. (2014) present a psychological perspective on the frequent failure of mandated disclosure, highlighting how limited attention can render complex disclosures an insufficient regulatory device. Recent large-scale field experiments by Shulman et al. (2024) provide compelling, real-world evidence that complexity directly undermines reader engagement: across millions of headline impressions in online news platforms, simpler linguistic constructions significantly increased click-through rates. This finding reinforces the idea that reducing textual complexity can meaningfully increase the likelihood that consumers will attend to, and process, disclosed information. Complementing evidence from controlled laboratory settings, Shulman et al. (2024) quantify the practical magnitude of a simplicity preference at population scale, showing that linguistic complexity alone measurably shifts consumer engagement in competitive attention markets.

Despite these advances, regulatory frameworks lack unified quantitative methods to evaluate how complexity and incentives jointly affect consumers' ability to effectively utilize disclosures in their decision-making. The closest paper to accomplishing this has been Colliard and Georg (2025), which develops an approach based on an algorithmic complexity measure first created by Halstead (1977). While their paper focuses on complex regulations from the perspective of those being regulated, our paper focuses on complex disclosure from the perspective of the consumers reading the disclosure. We similarly utilize experiments to measure complexity, however the present paper's manipulation of decision-making stakes allows us to more acutely measure cognitive costs and welfare implications of complex disclosures.

The broader economic literature extensively acknowledges decision complexity as a fundamental driver of decision quality and mistakes. Oprea (2020) demonstrates experimentally how increasing complexity significantly impairs decision quality through cognitive overload in strategic economic environments. Puri (2025) demonstrates how economic agents can have preferences for simplicity in lottery decisions due to cognitive constraints. Arrieta and Nielsen (2024) illustrate procedural decision-making responses to complexity, highlighting the qualitative effectiveness of structured choice procedures in complex decision problems, but noting a significant gap in quantifiable error and cost metrics.

A related strand of literature investigates the interplay of incentives and decision quality, particularly under conditions of complexity. Camerer and Hogarth (1999) comprehensively document how incentives typically enhance performance but exhibit diminishing or adverse returns as complexity escalates. Enke and Graeber (2023) empirically explore

how incentive structures interact with cognitive biases and errors under complex decisions.

The present paper utilizes the State-Dependent Stochastic Choice (SDSC) approach to measure cognitive economic costs from observed decision errors. Caplin and Dean (2015) introduced SDSC to explicitly quantify cognitive constraints through observed choice variability. Subsequent empirical validations by Caplin et al. (2020) demonstrate SDSC's capability to robustly quantify cognitive costs within rational inattention models. However, these powerful SDSC methodologies have not yet been systematically leveraged to explore how complexity and incentives jointly shape decision-making errors—precisely the empirical gap facing regulatory and marketing contexts.

Collectively, the economic and regulatory literatures clearly articulate the critical need for robust, integrated quantitative measures of consumer comprehension and decision quality under complexity. Our paper fills this critical gap by empirically quantifying how incentives and complexity jointly influence the structure of decision errors through the SDSC framework, providing a unified cognitive-economic metric relevant to regulators and marketers alike. By explicitly addressing this shared need, we offer actionable insights and methodological advancements that bridge theory and practice across economic theory, consumer psychology, regulatory disclosure practices, and applied marketing.

7 From Prototype to Policy Platform: Extending the Cognitive Economic Curve

The framework introduced in this paper offers more than a refined behavioral metric: it opens the door to a new science of disclosure evaluation. For decades, scholars and policymakers have lamented the inability to empirically distinguish between policies that are confusing and those that are clear, between disclosures that merely inform and those that truly empower. With cognitive economic curves, we now have a method that is not only grounded in economic theory but directly applicable in practice—portable across domains, adaptable to complexity, and aligned with the regulatory values of transparency, fairness, and effectiveness.

We are optimistic—and indeed, excited—about the future trajectory of this work. Below we sketch four natural extensions that will allow this framework to evolve from experimental prototype into a versatile platform for real-world evaluation and redesign.

Richer Decision Contexts: Beyond Accept/Reject While our initial application focuses on binary accept/reject decisions under varying complexity and incentives, many real-world policy contexts involve choice among multiple options or more nuanced tradeoffs. Choosing between insurance plans, benefit structures, or loan offers often involves multi-attribute comparisons. The cognitive economic curve can be extended to such settings by measuring accuracy not in terms of correctness per se, but in terms of alignment between predicted welfare and observed choices. Generalizations to multinomial or continuous choice settings—underpinned by state-dependent stochastic choice theory—are a natural next step Caplin et al. (2020).

Dynamic and Sequential Decisions Many disclosures are not one-shot interactions but part of longer decision episodes—filing taxes, enrolling in benefits, navigating consent interfaces, or managing digital privacy. Cognitive demands accumulate over time, and the challenge becomes not just clarity at a point in time but *navigability* over time. We envision dynamic extensions of cognitive economic curves that track evolving decision quality and attention costs across stages, using incentive-compatible path-dependent elicitation methods.

Targeted Design Through Population Heterogeneity The very same policy may be clear to one user and incomprehensible to another. Older adults, non-native English speakers, and low-literacy populations face systematically higher cognitive barriers. Cognitive economic curves can be estimated conditional on observable characteristics, enabling regulators to diagnose differential impacts and move beyond a “one-size-fits-all” model of communication. They can thus become tools not just for efficiency, but also for equity and inclusion.

Embedding in Policy Pilots and AI-Aided Redesign This method is ready to be embedded in the real policy design cycle. Disclosure templates can be pre-tested using this framework before rollout. Competing versions of benefit letters or digital consent flows can be benchmarked using empirical welfare curves. Importantly, this framework also aligns naturally with AI-aided redesign: generative language models can be tasked with producing candidate disclosures, which are then *scored* via cognitive economic experiments. The result is a virtuous cycle of design, testing, and improvement.

A First Application: Medicare Part D as a Benchmark Case Among the many settings in which cognitive economic curves could be deployed, Medicare Part D stands out as an

ideal early application. The policy is longstanding, the evidence of confusion is robust, and the welfare implications are large. Prior studies have shown that many enrollees fail to choose cost-minimizing plans, often due to complexity in how drug coverage, tiered copayments, and premium structures are presented (Abaluck and Gruber, 2011; Kling et al., 2012). Existing decision aids have had only modest success, and few tools allow for rigorous, welfare-based evaluation of proposed simplifications.

Our framework could be used to systematically compare competing plan interfaces, highlight which features most impede accurate decision-making, and measure whether simplified presentations meaningfully improve welfare. This would move the field beyond observational studies toward a structured, experimental method for improving choice architecture in one of the most consequential public programs in the United States.

We close not with finality, but with invitation. This method is simple enough to be deployed today, and rigorous enough to provide a new foundation for regulatory communication science. What has long been missing is now available: a tool to measure what clarity really means, and what it really delivers.

References

- Abaluck, Jason and Jonathan Gruber (2011) “Heterogeneity in choice inconsistencies among the elderly: Evidence from prescription drug plan choices,” *American Economic Review*, 101 (3), 377–381, 10.1257/aer.101.3.377.
- Allcott, Hunt and Judd B. Kessler (2019) “The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons,” *American Economic Journal: Applied Economics*, 11 (1), 236–276, 10.1257/app.20170328.
- Arrieta, Gonzalo and Kirby Nielsen (2024) “Procedural Decision-Making In The Face Of Complexity,” *Working Paper*, <https://garrieta.github.io/jmp/JMP.pdf>.
- Ben-Shahar, Omri and Carl E. Schneider (2014) *More Than You Wanted to Know: The Failure of Mandated Disclosure*, Princeton University Press, <https://www.jstor.org/stable/j.ctt5hhrqj>.
- Bhargava, Saurabh, George Loewenstein, and Justin Sydnor (2017) “Choose to Lose: Health Plan Choices from a Menu with Dominated Options,” *The Quarterly Journal of Economics*, 132 (3), 1319–1372, <https://www.jstor.org/stable/26372704>, Publisher: Oxford University Press.

- Camerer, Colin F. and Robin M. Hogarth (1999) "The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework," *Journal of Risk and Uncertainty*, 19, 7–42, <https://doi.org/10.1023/A:1007850605129>.
- Caplin, Andrew, Dániel Csaba, John Leahy, and Oded Nov (2020) "Rational Inattention, Competitive Supply, and Psychometrics," *The Quarterly Journal of Economics*, 135 (3), 1681–1724, <https://doi.org/10.1093/qje/qjaa011>.
- Caplin, Andrew and Mark Dean (2015) "Revealed Preference, Rational Inattention, and Costly Information Acquisition," *American Economic Review*, 105 (7), 2183–2203, <https://www.aeaweb.org/articles?id=10.1257/aer.20140117>.
- Chater, Nick and George Loewenstein (2023) "The i-frame and the s-frame: How focusing on individual-level solutions has led behavioral public policy astray," *Behavioral and Brain Sciences*, 46, e147, 10.1017/S0140525X22002023.
- Colliard, Jean-Edouard and Co-Pierre Georg (2025) "Measuring Regulatory Complexity," 10.2139/ssrn.3523824.
- Enke, Benjamin and Thomas Graeber (2023) "Cognitive uncertainty," *The Quarterly Journal of Economics*, 138 (4), 2021–2067, <https://doi.org/10.1093/qje/qjad025>.
- Financial Conduct Authority (2022a) "FCA Handbook – Principles for Businesses (PRIN) 2A.5," <https://www.handbook.fca.org.uk/handbook/PRIN/2A/5.html>.
- (2022b) "Final non-Handbook Guidance for firms on the Consumer Duty," Finalised Guidance FG22/5, London, UK, <https://www.fca.org.uk/publication/finalised-guidance/fg22-5.pdf>.
- (2022c) "A New Consumer Duty: Feedback to CP21/36 and Final Rules," Policy Statement PS22/9, London, UK, <https://www.fca.org.uk/publication/policy/ps22-9.pdf>.
- Flesch, Rudolph (1948) "A new readability yardstick," *Journal of Applied Psychology*, 32 (3), 221–233, 10.1037/h0057532.
- Florida Statutes (2024) "§ 627.4145. Readable language in insurance policies," <https://www.flsenate.gov/Laws/Statutes/2024/627.4145>.
- Frederick, Shane (2005) "Cognitive reflection and decision making," *Journal of Economic Perspectives*, 19 (4), 25–42, <http://dx.doi.org/10.1257/089533005775196732>.

- Gabaix, Xavier and David Laibson (2006) "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *The Quarterly Journal of Economics*, 121 (2), 505–540, <https://doi.org/10.1162/qjec.2006.121.2.505>.
- Halstead, Maurice H. (1977) *Elements of Software Science (Operating and programming systems series)*, USA: Elsevier Science Inc.
- Ippolito, Pauline M. (1986) "Consumer Protection Economics: A Selective Survey," in *Empirical Approaches to Consumer Protection Economics*, <https://www.ftc.gov/sites/default/files/documents/reports/empirical-approaches-consumer-protection-economics/198404consumereconomics.pdf>.
- Kling, Jeffrey R, Sendhil Mullainathan, Eldar Shafir, Lee Vermeulen, and Marian V Wrobel (2012) "Comparison friction: Experimental evidence from Medicare drug plans," *Quarterly Journal of Economics*, 127 (1), 199–235, [10.1093/qje/qjr055](https://doi.org/10.1093/qje/qjr055).
- Litman-Navarro, Kevin (2019) "We Read 150 Privacy Policies. They Were an Incomprehensible Disaster," *The New York Times*, <https://www.nytimes.com/interactive/2019/06/12/opinion/facebook-google-privacy-policies.html>.
- Loewenstein, George, Cass R. Sunstein, and Russell Golman (2014) "Disclosure: Psychology Changes Everything," *Annual Review of Economics*, 6, 391–419, <https://doi.org/10.1146/annurev-economics-080213-041341>.
- McDonald, Aleecia M. and Lorrie Faith Cranor (2008) "The Cost of Reading Privacy Policies," *I/S: A Journal of Law and Policy for the Information Society*, 4, 543, <https://heinonline.org/HOL/Page?handle=hein.journals/isjlp4&id=563&div=&collection=>.
- New York Insurance Law (2014) "§ 3102. Requirements for the use of readable and understandable insurance policies," <https://www.nysenate.gov/legislation/laws/ISC/3102>.
- Oprea, Ryan (2020) "What Makes a Rule Complex?" *American Economic Review*, 110 (12), 3913–51, <https://www.aeaweb.org/articles?id=10.1257/aer.20191717>.
- Puri, Indira (2025) "Simplicity and Risk," *The Journal of Finance*, 80 (2), 1029–1080, [10.1111/jofi.13417](https://doi.org/10.1111/jofi.13417), eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.13417>.

Shulman, Hillary C., David M. Markowitz, and Todd Rogers (2024) "Reading dies in complexity: Online news consumers prefer simple writing," *Science Advances*, 10 (23), eadn2555, 10.1126/sciadv.adn2555, Publisher: American Association for the Advancement of Science.

Solove, Daniel J. (2021) "The Myth of the Privacy Paradox," *SSRN Working Paper*, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3536265.

Weber, Ernst Heinrich (1834) *De pulsu, resorptione, auditu et tactu. Annotationes anatomicae et physiologicae, auctore.*, Leipzig: C.F. Koehler.

A Proofs

A.1 \bar{U} is non-decreasing

Theorem 2. *With NIAC, $\bar{U}(d)$ is non-decreasing in d .*

Proof. Take $d_2 > d_1 > 0$ and suppose to the contrary that

$$\frac{U(d_1)}{d_1} = \bar{U}(d_1) < \bar{U}(d_2) = \frac{U(d_2)}{d_2}.$$

In combination these yield

$$\begin{aligned} (d_2 - d_1) \left(\frac{U(d_1)}{d_1} - \frac{U(d_2)}{d_2} \right) &> 0 \\ \iff \frac{d_2}{d_1} U(d_1) + \frac{d_1}{d_2} U(d_2) &> U(d_1) + U(d_2). \end{aligned}$$

This contradicts NIAC since directly swapping the data between the two data sets allows the LHS sum to be achieved, which cannot be improving. This follows simply from using $a_{d_1}(n)$ in place of $a_{d_2}(n)$ when switching data set \mathbf{P}^{d_2} to A^{d_1} and conversely $a_{d_2}(n)$ in place of $a_{d_1}(n)$ when switching data set \mathbf{P}^{d_1} to A^{d_2} . \square

A.2 Proof of Theorem 1

The fact that $\bar{U}(d)$ is non-decreasing in d allows us to apply a wonderful general version of the Fundamental Theorem of Calculus: Rockafellar theorem 24.2.

Rockafellar, Theorem 24.2: Let $g : \mathbb{R}_+ \rightarrow \bar{\mathbb{R}}$ be a non-decreasing function with $g(0)$ finite. Let G be a closed convex function such that $G'_-(x) \leq g(x) \leq G'_+(x)$. Then,

$$G(x) = \int_0^x g(t) dt + \alpha$$

for some $\alpha \in \mathbb{R}$. Note that for a convex function to prove $G'_-(x) \leq g(x) \leq G'_+(x)$ involves a global proposition that the function lies above the supporting line at x with slope $g(x)$

$$G(y) \geq G(x) + (y - x)g(x).$$

Proof. For any $d > 0$, given state-dependent stochastic choice data, \mathbf{P}^d , define the corre-

sponding net utility,

$$\begin{aligned} N(d) &\equiv U(d) - K(d) \\ &= \sum_{\omega} \sum_n u(a_d(n), \omega) \mathbf{P}^d(n, \omega) - K(d) \end{aligned}$$

Consider an arbitrary $d_0 > 0$. This is associated with net utility $N(d_0)$ and normalized expected utility $\bar{U}(d_0)$. Consider now any other $d > 0$. Given optimality of data at d , \mathbf{P}^d must do at least as well as applying \mathbf{P}^{d_0} :

$$\begin{aligned} N(d) &\geq \sum_{\omega} \sum_n u(a_d(n), \omega) \mathbf{P}^{d_0}(n, \omega) - K(d_0) \\ &= d \frac{U(\mathbf{P}^{d_0})}{d_0} - K(d_0) \\ &= d\bar{U}(d_0) - d_0\bar{U}(d_0) + d_0\bar{U}(d_0) - K(d_0) \\ &= (d - d_0)\bar{U}(d_0) + N(d_0); \end{aligned}$$

where the second line follows from the definitions of $u(a_d(n), \omega)$ and $u(a_1(n), \omega)$, the third line involves adding and subtracting $d_0\bar{U}(d_0)$, and the final line applies the definitions of $N(d_0)$ and $\bar{U}(d_0)$. It follows immediately that N is convex and that $\bar{U}(d_0)$ is an element of the subdifferential of N at point d_0 ,

$$\bar{U}(d_0) \in \partial N(d_0).$$

Rockafellar Theorem 24.2 (above) applied to convex function N (closed for standard reasons of optimization) and non-decreasing function \bar{U} with $\bar{U}(d) \in \partial N(d)$ for all d yields,

$$N(d) = \int_0^d \bar{U}(t) dt + C,$$

for some constant $C \in \mathbb{R}$. Computing now the RHS of (1),

$$d\bar{U}(d) - \int_0^d \bar{U}(t) dt = d\bar{U}(d) - N(d) + C.$$

To pin down C , note that,

$$\lim_{d \searrow 0} N(d) = \lim_{d \searrow 0} \{U(P_d) - K(d)\} = - \lim_{d \searrow 0} K(d),$$

and

$$\lim_{d \searrow 0} N(d) = \lim_{d \searrow 0} \int_0^d \bar{U}(t) dt + C = C.$$

Hence, $C = - \lim_{d \searrow 0} K(d)$. But since the inattentive strategy is always available and costless in order to maximize welfare we need that $\lim_{d \searrow 0} K(P_d) = 0$. We conclude that $C = 0$. Applying in turn the definitions of $N(d) = U(d) - K(d)$ and $U(d) = d\bar{U}(d)$ we conclude,

$$\begin{aligned} - \int_0^d \bar{U}(t) dt &= -N(d) \\ &= K(d) - U(d) \\ &= K(d) - d\bar{U}(d), \end{aligned}$$

establishing (1) and completing the proof. □

B Additional Figures

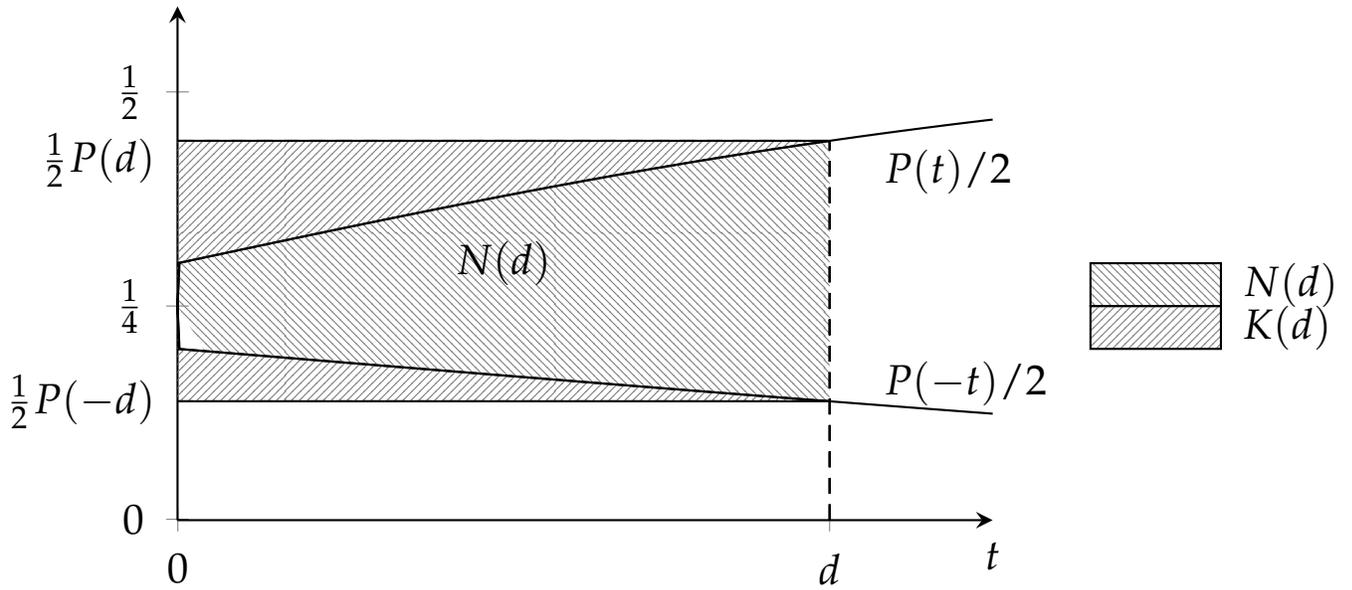


Figure 15: This “folded” version of the cognitive economic curve of Figure 2c illustrates, for $\mu = \frac{1}{2}$, how gross utility $U(d) = d[\mu P(d) - (1 - \mu)P(-d)]$ is the area of the rectangle; net utility is revealed to be the integral of $\bar{U}(t) = \mu P(t) - (1 - \mu)P(-t)$ and hence the area between the prior-weighted psychometric curve $\mu P(t)$ and its mirrored version $(1 - \mu)P(-t)$; and information costs $K(d) = U(d) - N(d)$ are consequently revealed by the rectangle’s area *not* between the curves.

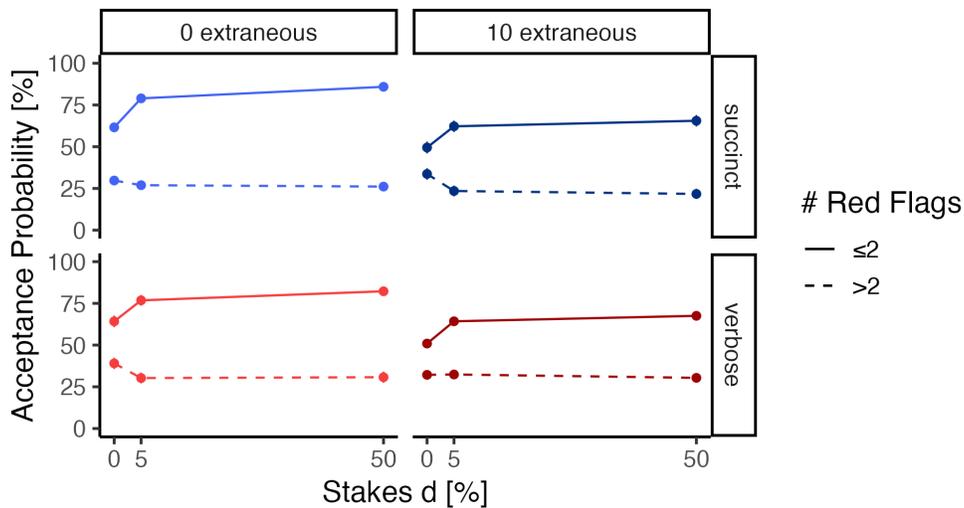


Figure 16: The “folded” cognitive economic curves, analogous to Figure 15, reveal the net welfare at each stake as the area between the curves up to that point.

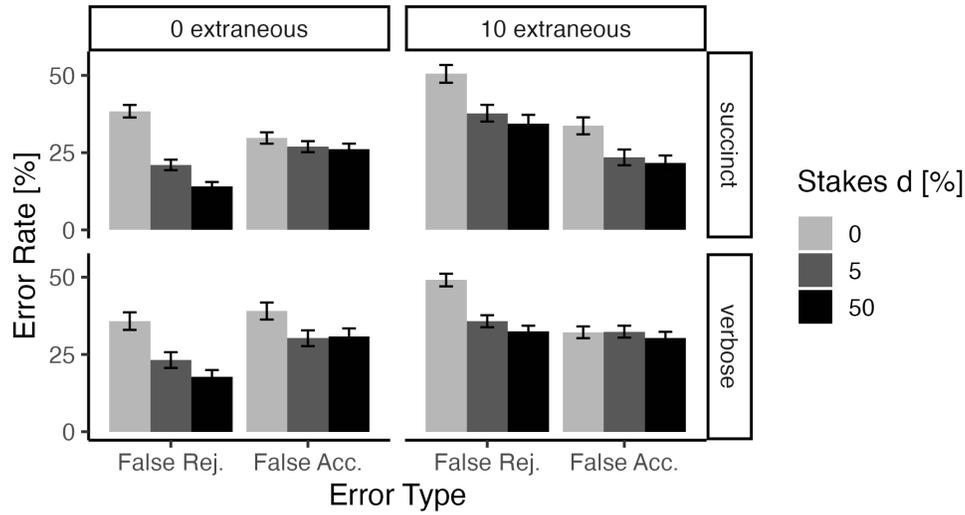


Figure 17: Error types by stakes.

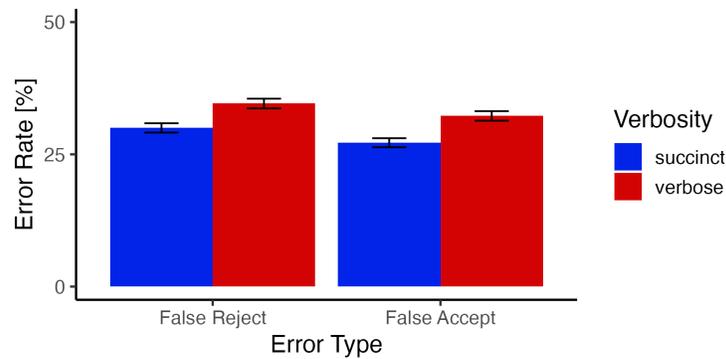


Figure 18: Error rates by verbosity. This figure suggests that verbosity increases both False Rejects and False Accepts, but this is an artifact of the experimental design's correlation between verbosity and the number of extraneous items (the experiment featured twice as many trials for the *succinct*×0 and *verbose*×10 treatments as for the *succinct*×10 and *verbose*×0 treatments). In Figure 13, False Rejects did not differ across verbosity levels when controlling for the number of extraneous items. The difference in False Rejects observed here is due to the fact that for *verbose* there are twice as many trials with 10 extraneous items than with 0 extraneous items, whereas for *succinct* there are twice as many trials with 0 extraneous items than with 10 extraneous items.

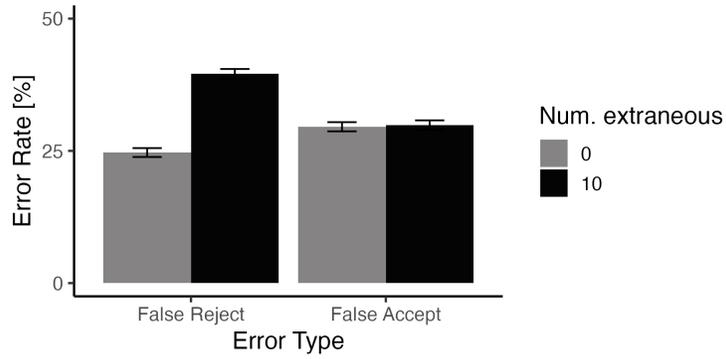


Figure 19: Error rates by number of extraneous items.

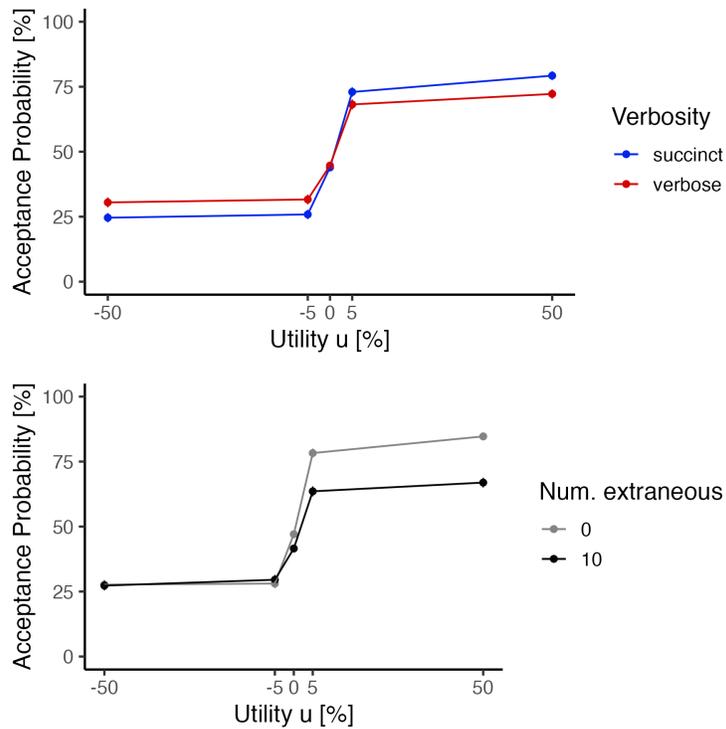


Figure 20: Cognitive economic curves, split by each treatment dimension separately. In interpreting these plots, the same caveat applies with regards to the correlated experiment design as in Figure 18.

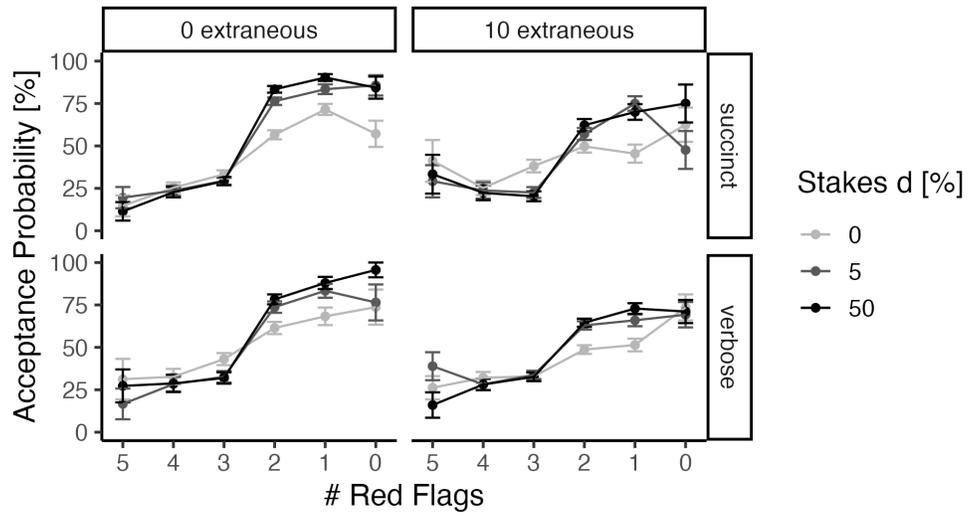


Figure 21: Psychometric curves by treatment and stakes. The curves steepen as the incentive increases.

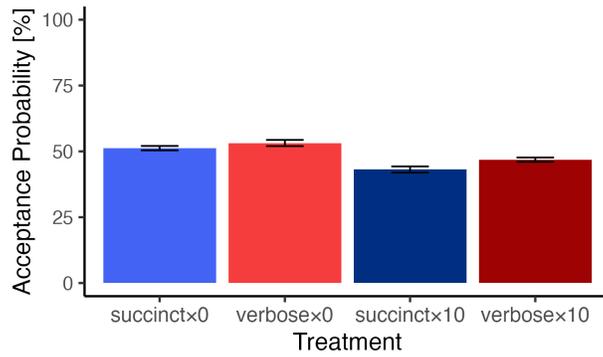


Figure 22: Acceptance probabilities by treatment.

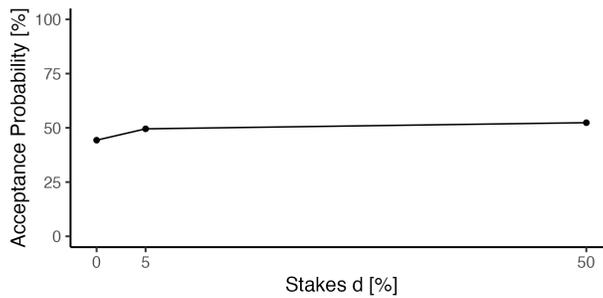


Figure 23: Acceptance probabilities as a function of stakes.

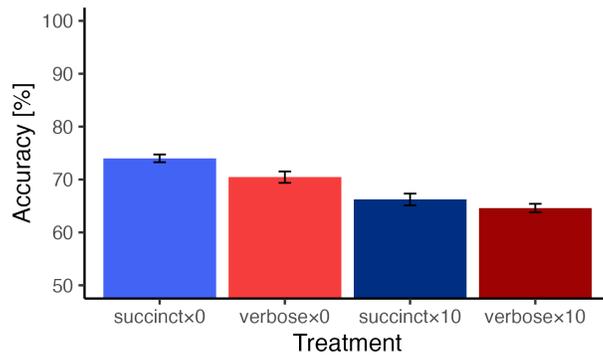


Figure 24: Overall accuracy by treatment.

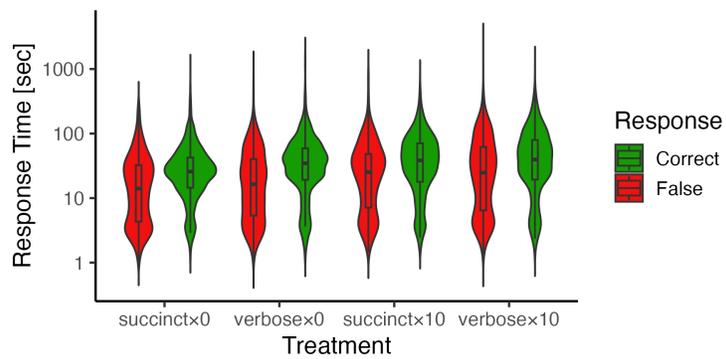


Figure 25: Response time by correctness.

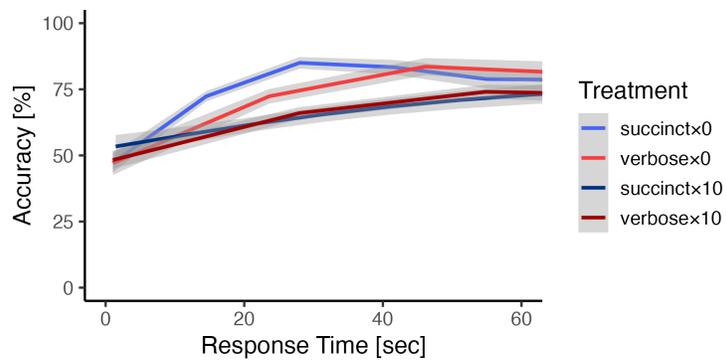


Figure 26: Accuracy as a function of response time.

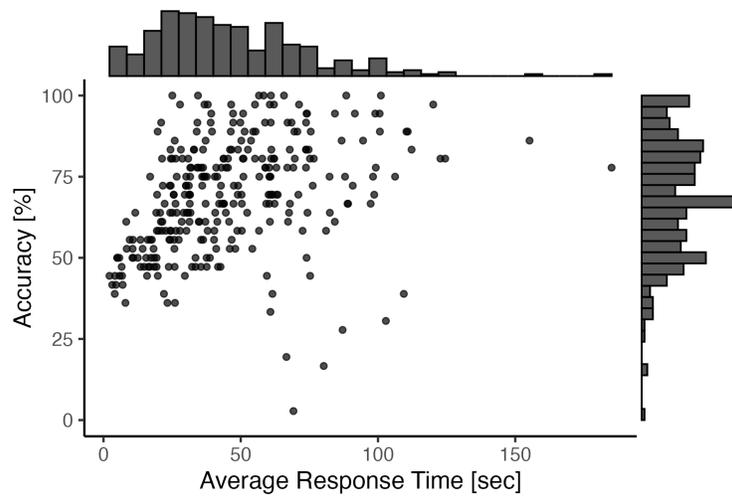


Figure 27: Scatterplot of each subject’s average accuracy and response times, which are positively correlated across subjects.

C Experiment timeline

C.1 Start Screen

Welcome to this experiment!

You will receive \$5 base pay for completing this experiment. You can earn up to an additional \$8, for a total of up to \$13.

This experiment will take around 35 minutes.

Next

Consent Screen

C.2 Instructions

Objective

You can win a \$8 bonus in this experiment, in addition to your base pay.

In each round of the experiment you **start with 50 coins**, and you will participate in a task in which you can potentially **add or subtract coins** to this. The more coins you end the task with, the higher your likelihood of earning the \$8 bonus payment.

At the end of the experiment, the computer will randomly select one round to determine your bonus. **The number of coins you collected in that round determines your chance of winning the \$8 bonus. The more coins you collect in each round, the higher your chance of receiving the \$8 bonus.**

[Next](#)

Task Description: Page 1 of 2

In each round, you will be shown a hypothetical **privacy policy** for a fictional product. These policies, like the ones you encounter in real life, describe a number of privacy and data security features of a fictional website or app.

The number of coins you earn in the round depend on whether you decide to accept or reject the fictional product based on the information contained in its privacy policy. Here is an example of the kind of policy you will be shown:

LagoonNet Privacy Policy

Last Updated: October 23, 2024

[The content of this privacy policy is intentionally blurred in the image.]

Example of a Hypothetical Privacy Policy (Contents will be unblurred in the actual experiment)

Click next to proceed to the next page.

Next

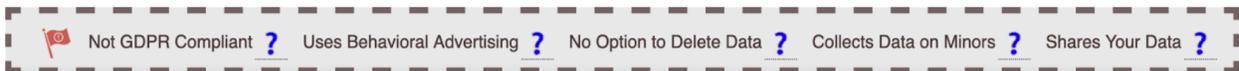
Task Description: Page 2 of 2

In addition to seeing a privacy policy, in each round you will also be a set of 5 "red flag" features, which are features of the product you should **avoid**. The privacy policy which tell you whether the product has these features.

Your task is to choose whether to **accept or reject** the product, given the information in its policy about red flag features. Your decision will determine whether you gain coins (relative to the 50 coins you start the round with) or lose coins.

- If you **ACCEPT** a product with **0, 1, or 2** red flag features , you will **gain** coins.
- If you **ACCEPT** a product with **3, 4, or 5** red flag features , you will **lose** coins.
- If you **REJECT** the product, you will **not win or lose any coins**.

Below us an **example** of how the red-flag features for the round are described to you:



In the experiment, you can hover your mouse over the blue  to see a short explanation of each red-flag feature. You can then read the privacy policy to determine whether each red-flag feature is present in the disclosure describing the fictional product.

For example, suppose you are told one "red flag" feature is "the website watches you through your webcam".

- The *red-flag* feature is **present** if the privacy policy explicitly mentions that the website uses your webcam to watch you.
- The *red-flag* feature is **not present** if the policy does not mention this, or explicitly mentions that the website **does not** use your webcam to watch you.

Click Next to proceed to the Payoff Rule.

[Next](#)

C.3 Understanding Check

Understanding Check

Fill out the following scenarios with the potential outcomes for when you accept or reject. Remember, you **start with 50 coins** at the beginning of each round!

You must answer all questions correctly to proceed. [Review task instructions](#)

Scenario A: If there are **50 coins at stake**:

| # of red flags  present in this policy | Potential outcome if you ACCEPT this policy | Potential outcome if you REJECT policy |
|---|--|---|
| 0, 1, 2 | <input type="text"/> | <input type="text"/> |
| 3, 4, 5 | <input type="text"/> | <input type="text"/> |

Scenario B: If there are **5 coins at stake**:

| # of red flags  present in this policy | Potential outcome if you ACCEPT this policy | Potential outcome if you REJECT policy |
|---|--|---|
| 0, 1, 2 | <input type="text"/> | <input type="text"/> |
| 3, 4, 5 | <input type="text"/> | <input type="text"/> |

Scenario C: If there are **0 coins at stake**:

| # of red flags  present in this policy | Potential outcome if you ACCEPT this policy | Potential outcome if you REJECT policy |
|---|--|---|
| 0, 1, 2 | <input type="text"/> | <input type="text"/> |
| 3, 4, 5 | <input type="text"/> | <input type="text"/> |

Next

- Clicking **Review task instructions** would reveal a scrollable box containing a replica of content from the **Instructions** screens.
- The drop-down menu in each answerable cell listed all possible outcomes (0 coins, 45 coins, 50 coins, 55 coins, 100 coins).
- Participants were required to select an answer for all twelve questions before submitting their answers by clicking **Next**. If all questions were answered correctly, participants proceeded directly to the **Task Start Screen**. Otherwise, participants were shown a notice listing incorrectly answered cells and were allowed to re-attempt the Understanding Check.
- After their second submission, participants who answered at least 10/12 questions correctly were permitted unlimited additional attempts at the Understanding Check until they correctly answered for all cells.
- Participants who failed to answer at least 10/12 questions correctly within two attempts were automatically redirected to a failure screen with no option to proceed further:

Notice: Failed Understanding Check

Unfortunately, you have failed to meet the comprehension criteria required to proceed with this experiment.

Please return your submission by closing this page and clicking **Stop Without Completing** on Prolific.

C.4 Task Start Screen

Start the task!

Congratulations! You have successfully completed the understanding check.

Click next when you are ready to begin the experiment.

Next

C.5 Choice Task (36 rounds)

Round 1: **50 coins at stake!**

[Proceed to the task](#)

Round 1: **50 coins at stake!**

Collects Location Information [?](#) Uses Cookies [?](#) Profiles User Behavior [?](#) Collects Device Information [?](#) Lacks Data Security Measures [?](#)

AtomicFlow Privacy Policy

Last Updated: May 01, 2025

We do not use your behavior and information to build comprehensive user profiles.

Information about you and your visits is not remembered by our company using cookies.

We obtain data about the particular device you use, such as your phone or computer.

We lack robust measures to secure your data against breaches or misuse.

We use your device's GPS or other location services to track your whereabouts.

Neither gain nor lose any coins →

REJECT!

ACCEPT!

← Gain 50 coins if 0, 1, or 2 red flag features
← Lose 50 coins if 3, 4, or 5 red flag features

You must scroll to the end of the policy to proceed.

- After participants scrolled to the bottom of the private policy, the **ACCEPT!** and **REJECT!** boxes became selectable:

Neither gain nor lose any coins →

REJECT!

ACCEPT!

← Gain 50 coins if 0, 1, or 2 red flag features
← Lose 50 coins if 3, 4, or 5 red flag features

[Submit my choice](#)

- Participants were given unlimited time to switch between options before confirming their selection by clicking **Submit my choice**.

C.6 Task End Screen

You are done!

You have completed all rounds.

Next, you will be asked to fill out a short survey. Press next to continue.

Next

C.7 Survey

- On each survey page, participants were required to provide an answer for all listed questions before they could proceed to the next screen.
- All open-response answer cells exclusively accepted numerical values.

Survey: Page 1 of 3

Please answer the following questions.

GalaxySoft Privacy Policy

Last Updated: March 25, 2025

The Company expressly states that it does not use user data for artificial intelligence system development. This policy of non-utilization encompasses all AI training activities. The Company has established procedures to prevent such data usage. The Company hereby affirms and represents that it does not retain user information indefinitely. This policy establishes defined retention periods and deletion schedules. The Company maintains strict protocols governing data deletion timelines.

The Company hereby notifies users that it disregards browser-based 'Do Not Track' signals. This encompasses continued tracking activities regardless of user preferences. The non-implementation of such signals follows Company procedures. The Company expressly states that it does not obtain additional user information from external entities. This policy of non-collection encompasses all third-party data sources. The Company has established procedures preventing external data acquisition. Pursuant to Company policy, we do not gather or maintain unique biometric measurements. This includes refraining from collecting fingerprints and facial recognition data. The Company has implemented measures preventing such data collection.

For **long policies in paragraphs** (as pictured above), what proportion of the time do you think you responded correctly (i.e., accepted when there were 0, 1, or 2 red flags, and rejected when there were 3, 4, or 5 red flags)?



AlphaCore Privacy Policy

Last Updated: March 24, 2025

Sensitive details like your race, religion, and sexual orientation are not collected by our company.

We collect data from individuals who have not yet reached 18 years of age.

For our European users, European data protection laws (GDPR) are completely adhered to by our company.

Detailed user profiles are created by our company using behavior and information.

You cannot easily remove all your information from our systems.

For **short policies in sentences** (as pictured above), what proportion of the time do you think you responded correctly (i.e., accepted when there were 0, 1, or 2 red flags, and rejected when there were 3, 4, or 5 red flags)?



Next

Survey: Page 2 of 3

Please answer the following questions.

Think about the requests you've seen on real websites asking whether or not they can track your browsing. When asked if you consent to tracking your browsing, do you usually:

- Accept all trackers
- Accept necessary trackers and some, but not all, optional trackers
- Accept only necessary trackers
- Reject all trackers, even if it means not using the website
- Reject all trackers only if you can continue using the website
- It depends on the website and the trackers used
- I don't know

In your everyday life, how long do you spend reviewing a typical privacy policy for a website or service you are interested in using? (Please answer in seconds)

[Next](#)

Survey: Page 3 of 3

Please answer the following questions.

What is your total annual income?

- \$0-\$30,000
- \$31,000-\$60,000
- \$61,000-\$90,000
- \$91,000-\$120,000
- \$120,000+
- Prefer not to say

What is the highest degree or level of school you have completed?

A bat and a ball cost \$1.10 in total. The bat costs \$1 more than the ball. How much does the ball cost? (Please answer in dollars)

If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets? (Please answer in minutes)

In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half the lake? (Please answer in days)

[Next](#)

- The drop-down menu for *What is the highest degree or level of school you have completed?* consisted of options: 'No schooling completed', 'Some high school, no diploma', 'High school graduate', 'Some college credit, no degree', 'Trade/technical/vocational training', 'Associate degree', 'Bachelor's degree', 'Master's degree', 'Professional degree', 'Doctorate degree', and 'Prefer not to say'.

Survey Complete!

You have completed the survey.

Proceed to the following page to determine your payoff. Press next to continue.

[Next](#)

- Note: Additional demographics information not collected in this survey (Age, Gender) were retrieved directly from Prolific.

C.8 Bonus Determination

The Raffle!

Now, all of your coins will be used to buy tickets for a raffle! If you hold the ticket with the winning number, you win the \$8 bonus! There are 100 total raffle tickets. Each raffle ticket costs one coin.

Round 28 had been selected for payoff! In this round, you stored **50 total coins**. This means you will buy 50 raffle tickets.

[BUY RAFFLE TICKETS!](#)

- For each participant, one round (of 36 total) would be randomly selected and the number N of stored coins for that round (0, 45, 50, 55, or 100) was recalled.
- N numbers were randomly drawn without replacement from the set $\{0, 1, 2, \dots, 99\}$ of 100 numbers to form a consideration set for the bonus determination. All selected numbers were revealed to the participant on the following screen after they clicked [BUY RAFFLE TICKETS!](#).

The Raffle!

Here are all the raffle tickets, with their corresponding numbers printed. **You bought the tickets printed in red!**

Raffle Tickets

| | | | | | | | | | |
|---|----|----|----|----|----|----|----|----|----|
| 9 | 19 | 29 | 39 | 49 | 59 | 69 | 79 | 89 | 99 |
| 8 | 18 | 28 | 38 | 48 | 58 | 68 | 78 | 88 | 98 |
| 7 | 17 | 27 | 37 | 47 | 57 | 67 | 77 | 87 | 97 |
| 6 | 16 | 26 | 36 | 46 | 56 | 66 | 76 | 86 | 96 |
| 5 | 15 | 25 | 35 | 45 | 55 | 65 | 75 | 85 | 95 |
| 4 | 14 | 24 | 34 | 44 | 54 | 64 | 74 | 84 | 94 |
| 3 | 13 | 23 | 33 | 43 | 53 | 63 | 73 | 83 | 93 |
| 2 | 12 | 22 | 32 | 42 | 52 | 62 | 72 | 82 | 92 |
| 1 | 11 | 21 | 31 | 41 | 51 | 61 | 71 | 81 | 91 |
| 0 | 10 | 20 | 30 | 40 | 50 | 60 | 70 | 80 | 90 |

You are holding 50 out of 100 tickets.

To determine your payoff, draw the winning number by stopping the clock below! You will win the \$8 raffle if you have collected the raffle ticket with the winning number. The winning number is the last two digits in the millisecond unit at which you stop the clock.

10:15:26.1??

STOP THE CLOCK!

- The N values in the consideration set were visualized as red tickets.
- The numbers on the clock continuously updated on the screen. When participants clicked **STOP THE CLOCK!**, the clock froze and the obscured millisecond digits (??) were revealed in place.

- If the last two digits of the stopped clock matched one of their ticket numbers, participants received the bonus payment:

10:15:26.198

Congratulations! Last two digits, 98, matched your tickets.

Next

- If the last two digits of the stopped clock did not match any ticket numbers, participants did not receive the bonus payment:

10:15:26.189

Unfortunately, last two digits, 89, are not held.

Next

C.9 End Screen

Results

You have completed the experiment!

The raffle determined that your total bonus is \$8.00 for a total payment of \$13.00.

Please click below to finish the experiment and return to Prolific to collect your payment.

Finish

D Choice Screen Examples

The following are selected example privacy policies under each treatment for the following **red flags**.¹¹

- **Collects Location Information:** The company tracks where you are or where you've been using your device's GPS or other location services.
- **Uses Cookies:** The company uses small files called cookies to remember information about you and your visits.
- **Profiles User Behavior:** The company creates detailed profiles about users based on their behavior and information.
- **Collects Device Information:** The company collects information about your specific device (like your phone or computer).
- **Lacks Data Security Measures:** The company doesn't have strong safeguards in place to protect your information from breaches or misuse.

D.1 Example: Succinct × 0 extraneous attributes

AtomicFlow Privacy Policy

Last Updated: May 01, 2025

We do not use your behavior and information to build comprehensive user profiles.

Information about you and your visits is not remembered by our company using cookies

We obtain data about the particular device you use, such as your phone or computer.

We lack robust measures to secure your data against breaches or misuse.

We use your device's GPS or other location services to track your whereabouts.

¹¹The full set of possible policy attributes and potential statements are in <https://osf.io/nhy8g>.

D.2 Example: Succinct × 10 extraneous attributes

AtomicFlow Privacy Policy

Last Updated: May 01, 2025

Information about you and your visits is not remembered by our company using cookies.

Sensitive details like your race, religion, and sexual orientation are collected by our company.

Other businesses cannot purchase your contact information from our company.

We do not collect data from individuals who have not yet reached 18 years of age.

Artificial intelligence systems are taught using your data by our company.

We do not use your behavior and information to build comprehensive user profiles.

We obtain data about the particular device you use, such as your phone or computer.

Personal details such as your age, gender, and income are gathered by our company

For our European users, our practices do not fully align with European data protection laws (GDPR).

We use your device's GPS or other location services to track your whereabouts.

We lack robust measures to secure your data against breaches or misuse.

We show you targeted ads based on information about your online activities.

We obtain data about you from your social media accounts.

For our California users, our practices do not fully align with California's privacy laws (CCPA).

We disregard the 'Do Not Track' preferences set in your browser.

D.3 Example: Verbose × 0 extraneous attributes

AtomicFlow Privacy Policy

Last Updated: May 01, 2025

The Company expressly states that it does not develop detailed user profiles from behavioral patterns or submitted information. This policy of non-profiling encompasses all user activities and data. The Company has established procedures to prevent such profile creation. The Company expressly states that it does not employ cookie technology in any form. This encompasses non-collection of user activity and site visitation data through cookies. The Company has established procedures to prevent cookie implementation.

The Company hereby notifies users that it gathers specific information about their devices, including computers and mobile equipment. This encompasses collection of device identifiers and specifications. The processing of such information follows established procedures. The Company hereby notifies users of the absence of robust information security measures. This encompasses insufficient safeguards against unauthorized access and misuse. The Company operates without comprehensive data protection protocols. The Company hereby notifies users that it systematically collects and processes location data through authorized tracking technologies. This includes information gathered via device GPS and other location services. The collection and retention of such geographic data occurs in accordance with Company procedures.

D.4 Example: Verbose × 10 extraneous attributes

AtomicFlow Privacy Policy

Last Updated: May 01, 2025

The Company expressly states that it does not employ cookie technology in any form. This encompasses non-collection of user activity and site visitation data through cookies. The Company has established procedures to prevent cookie implementation. The Company hereby notifies users of its collection and retention of sensitive personal information through authorized channels. This encompasses data including, but is not limited to, information about race, religious affiliation, and sexual orientation. The processing of such sensitive information is governed by internal data handling procedures.

The Company expressly states that it does not sell user contact information to external businesses. This policy of non-sale encompasses all contact detail transfers. The Company has established procedures to prevent such commercial transactions. The Company expressly states that it does not collect information from individuals under eighteen (18) years of age. This policy of non-collection encompasses all minor user data. The Company has established procedures to prevent gathering of such information. The Company hereby notifies users that their data is utilized in artificial intelligence system training. This encompasses the processing of user information for AI development purposes. The implementation of such training follows established protocols.

The Company expressly states that it does not develop detailed user profiles from behavioral patterns or submitted information. This policy of non-profiling encompasses all user activities and data. The Company has established procedures to prevent such profile creation. The Company hereby notifies users that it gathers specific information about their devices, including computers and mobile equipment. This encompasses collection of device identifiers and specifications. The processing of such information follows established procedures. The Company hereby notifies users that it systematically collects and maintains demographic information through authorized channels. Such demographic information includes, but is not limited to, user age, gender, and income-related data. This information is processed and retained in accordance with established Company protocols.

The Company hereby notifies European users of incomplete compliance with GDPR requirements. This encompasses various aspects of data protection regulations. The Company's practices do not fully satisfy GDPR standards. The Company hereby notifies users that it systematically collects and processes location data through authorized tracking technologies. This includes information gathered via device GPS and other location services. The collection and retention of such geographic data occurs in accordance with Company procedures. The Company hereby notifies users of the absence of robust information security measures. This encompasses insufficient safeguards against unauthorized access and misuse. The Company operates without comprehensive data protection protocols.

The Company hereby notifies users that it processes online behavioral data to deliver targeted advertisements. This encompasses the analysis of user online activities for promotional purposes. The implementation of such targeted advertising follows established Company procedures. The Company hereby notifies users of its collection of information from social media platforms and services. This encompasses data gathering from social networking accounts and activities. The processing of such social media information follows established procedures.

The Company hereby notifies California users of incomplete compliance with CCPA requirements. This encompasses various aspects of privacy regulations. The Company's practices do not fully satisfy CCPA standards. The Company hereby notifies users that it disregards browser-based 'Do Not Track' signals. This encompasses continued tracking activities regardless of user preferences. The non-implementation of such signals follows Company procedures.