

RATIONAL INATTENTION, COMPETITIVE SUPPLY, AND PSYCHOMETRICS*

Andrew Caplin

Dániel Csaba

John Leahy

Oded Nov

April 6, 2020

Abstract

We introduce a simple method of recovering attention costs from choice data. Our method rests on a precise analogy with production theory. Costs of attention determine consumer demand and consumer welfare just as a competitive firm's technology determines its supply curve and profits. We implement our recovery method experimentally, outline applications, and link our work to the broader literature on inattention and mistaken decisions.

JEL Codes: D83, D90.

*We thank Ernst Fehr for guiding us in the direction of simplification before we knew there was anything simple to be found. We thank Sandro Ambuehl, Roland Bénabou, Mark Dean, Sam Gershman, Andrei Gomberg, Wei Ji Ma, Daniel Martin, David Redish, Joyce Sadka, Isabelle Salcher, Eric Spurlino, and Dmitry Taubinsky for helpful discussions and three anonymous referees for constructive comments and suggestions. We thank Han Su and Theodore Kim for programming the experimental interface. We thank the Alfred P. Sloan and NOMIS Foundations for support. Corresponding author: Andrew Caplin, email: andrew.caplin@nyu.edu, tel: (+1) 212 998 8950, 19 W. 4th Street, 6th Floor, New York, NY 10012. Word count: 18,319.

1 Introduction

A defining feature of the modern era is our access to effectively unlimited amounts of information. Yet scarcity of inputs such as time and mental effort, and limits on our ability to learn from these inputs, make it effectively impossible to know everything that might be of value. The resulting constraints on information processing determine consumer demand and consumer welfare just as a competitive firm’s technology determines its supply curve and its profitability. Important as these constraints on information processing might be, it has proven challenging to assess their role due to their essentially subjective nature.

In this paper we adapt standard methods from production theory to the analysis of attention and comprehension. We show precisely how one can recover costs of attention from suitably rich data on choices. This bridges a key gap between theory and practice by allowing the role that costs of attention play in generating decision making errors to be studied systematically.

Our method rests on making the analogy between a consumer’s costs of attention and a competitive firm’s costs of production precise. There are two key links. First, we linearly scale up and down attentional incentives just as the price faced by a competitive firm linearly scales its production incentive. Second, we normalize utility by removing the confounding effect of incentives. This is analogous to measuring the output of a firm rather than its profit. Our incentive-based psychometric curve plots the linear incentive against normalized utility, just as the competitive firm’s supply curve plots its chosen output level against the price of output. Costs of attention are recovered from this curve just as a competitive firm’s total costs of production are recovered from its marginal cost curve, as reflected in output choices at different prices. Likewise we can identify consumer welfare net of costs precisely as one can identify producer surplus from the supply curve of a competitive firm. We provide tests for the existence of any rationalizing cost function analogous to revealed profitability tests [Varian, 2014]. We also develop a “revealed more costly than” binary relation that extracts rich information from each experimental observation.

Our recovery method uses state dependent stochastic choice (SDSC) data. This records the impact on choice probabilities of changes in measurable features of the environment, such as tax rates or product attributes. Since recovery is based on marginal logic, it also requires variation in incentives.

To illustrate our methods we gather essentially ideal data in a purpose-built psychometric experiment that is closely related to simple production tasks that have been analyzed in the field (e.g. [Guiteras and Jack \[2018\]](#), [Kaur et al. \[2019\]](#)). We confirm that the resulting experimental data satisfy simple behavioral tests for the existence of an attention cost function that rationalizes the data. We estimate cost functions corresponding to different levels of task difficulty and compute corresponding levels of welfare.

How well attentionally-demanding activities are performed in the absence of incentives is task specific. We identify it by running experimental treatments with incentives set to zero for each level of task difficulty. This allows us to recover not only the incentive-based psychometric curve, but also the classical psychometric curve in which all that varies is task complexity (as in classical experiments dating back to [Weber \[1834\]](#) and [Fechner \[1860\]](#)). By varying both the incentive and stimulus strength we produce a psychometric surface that quantifies the trade-off between task complexity and incentives.

Just as economists analyze production costs as based on the inputs of labor, raw materials, and capital, so psychologists analyze subjectively costly inputs that determine cognition, including processing time, intensity of effort, and energy use. Psychological research on attention focuses in particular on the drift-diffusion framework, in which the process of learning plays out over time [[Ratcliff, 1978](#), [Fehr and Rangel, 2011](#)]. In fitting with the process-based approach, we find that estimated attention costs in our psychometric task are highly correlated with decision time. As has been the case with costs of production, input-based decompositions of the revealed costs of attention may make cost estimates portable between distinct attentional challenges that call on common resources. There is also an interesting interplay between incentives and time. In the psychometric task there is a non-trivial time input even without a reward. This is task specific and suggests that the disutility (or utility) of attention may be important to uncover.

Our work contributes to a larger body of recent applied and theoretical work on clearly low quality decisions. Systematic decision making errors have been identified in case after case, be it failure to internalize sales taxes, [[Chetty et al., 2009](#), [Taubinsky and Rees-Jones, 2018](#), [Morrison and Taubinsky, 2019](#)], incomplete take up of the Earned Income Tax Credit, [[Currie, 2008](#), [Bhargava and Manoli, 2015](#)], or in choices related to medical insurance and medication [[Bhargava et al., 2017](#), [Handel and Kolstad, 2015](#), [Abaluck et al., 2018](#), [Kling et al., 2012](#)]. Changes in presentation and user interface are often

found to influence the quality of decisions suggesting that the difficulty in comprehending information on which we focus is a key piece of the puzzle.

While costs are part of the explanation for low quality decisions, there are many other plausible explanations that have been put forward for these disparate phenomena. [Handel and Schwartzstein \[2018\]](#) make a valuable distinction between forces that are predominantly “frictional” related to costly information acquisition, including search [[Stigler, 1961](#)], rational inattention [[Sims, 1998](#)], and the drift-diffusion model [[Fehr and Rangel, 2011](#)], and forces related to “mental gaps” that involve deeper departures from standard rational choice analysis. One form of mental gap that is outside our rational inattention framework arises when utility weights of distinct attributes of available goods (e.g. price and quality) depend on the nature of the choice set itself. In particular, salience theory is based on the idea that characteristics that “stand out” attentionally may be over-weighted [[Bordalo et al., 2012](#), [Taylor and Thompson, 1982](#)]. There are also essentially automatic attentional phenomena that rational inattention theory has nothing to say about and yet which may in many cases be dominant [[Dehaene et al., 2017](#)]. Past experiences and memories that choice situations trigger may shape initial beliefs, perceived utility and also limit what is learned [[Malmendier and Nagel, 2011](#), [Bordalo et al., 2020](#), [Wachter and Kahana, 2019](#)]. Another form of mental gap that economists have studied is selective attention introduced in [[Schwartzstein, 2014](#)]. For example [Hanna et al. \[2014\]](#) study seaweed farmers in Indonesia who never even observe let alone learn the impact of one crucial determinant of productivity (the size of the “pod”). Logical fallacies concerning sources of information have also been identified, including correlation neglect [[Akerlof and Shiller, 2015](#), [Enke and Zimmermann, 2017](#)] and neglect of clear rules for the selection of new information (e.g. in a partisan news source) that should rationally impact interpretation [[Enke, 2020](#)]. In cases such as these in which priors may be hard to rationalize, it becomes important to find methods for accurately measuring them along with any relevant costs of learning. Hybrid approaches allowing for mixtures of forces are a high priority going forward.

In section 2 we introduce the rational inattention (RI) model and link it to applications. We also introduce the incentive-based psychometric curve (IPC) and the connection with competitive supply. In section 3 we establish the recovery theorem and relate it to earlier recovery results. In section 4 we illustrate the value of linking attention theory to elementary microeconomics to test the theory of

profit maximization. Experimental design is in section 5. The experimental results are in section 6. Section 7 links to applications. Section 8 concludes. The Appendices include experimental instructions and experimental robustness checks.

2 The Incentive-based Psychometric Curve (IPC)

We now introduce the formal language required for our recovery result and connect it with applications. We model decision making environments in which the optimal choice is initially uncertain, and in which it takes attentional effort to improve understanding enough to choose appropriately. For example, choice might involve selecting one of a large set of available medical insurance plans. The decision maker is initially uncertain about which is the best plan given their particular medical history and prognosis. Between the insurance plans and appropriate medical experts there is in principle rich enough information on medical history and on coverage restrictions, deductibles, co-pays, and in network care facilities, to determine the plan that maximizes expected utility. Yet few are dedicated or expert enough to fully assess all plans. The end result is that there are widespread mistakes, with clearly dominated options chosen in many cases [Bhargava et al., 2017, Handel and Kolstad, 2015, Abaluck et al., 2018, Kling et al., 2012].

If accurately internalizing all potentially available information was costless, there would be far fewer mistakes made. The fact that there are so many mistakes suggests to the contrary that learning is costly. RI theory models investment in attention that optimally balances the subjectively assessed costs of learning against the subjectively assessed improvement in utility that such learning brings about. In the medical case, what is learned will likely depend on details of the individual's medical history. Someone with diabetes is likely to be more attuned to information related to costs of insulin, while someone with children might focus more on costs of pediatric care. How well informed is the final decision depends on how much of the potentially available information they access, and how well they are able to comprehend this information. To take a common-place example, those who are unfamiliar with computer programming are unlikely even to read about potential fixes to buggy software. In addition to the rationalizable frictional mistakes that we model, other attentional forces related to mental gaps are often in play related e.g. to which features of the contract are most salient and the

impact this has on attribute weights [Bordalo et al., 2012].

2.1 Model and Applications

Figure 1 presents the motivation for our approach. Figure 1a is a classical supply curve which provides a geometric representation of total revenue, total cost and producer surplus for a competitive firm. Next to it in figure 1b is the IPC, relating incentives that mimic prices and “attentional output” that mimics firm output, with the axes flipped. By analogy with competitive supply the recovery theorem will allow us to identify the equivalents of total revenue, total cost and producer surplus for the attentional problem. In this section we define the axes and the IPC.

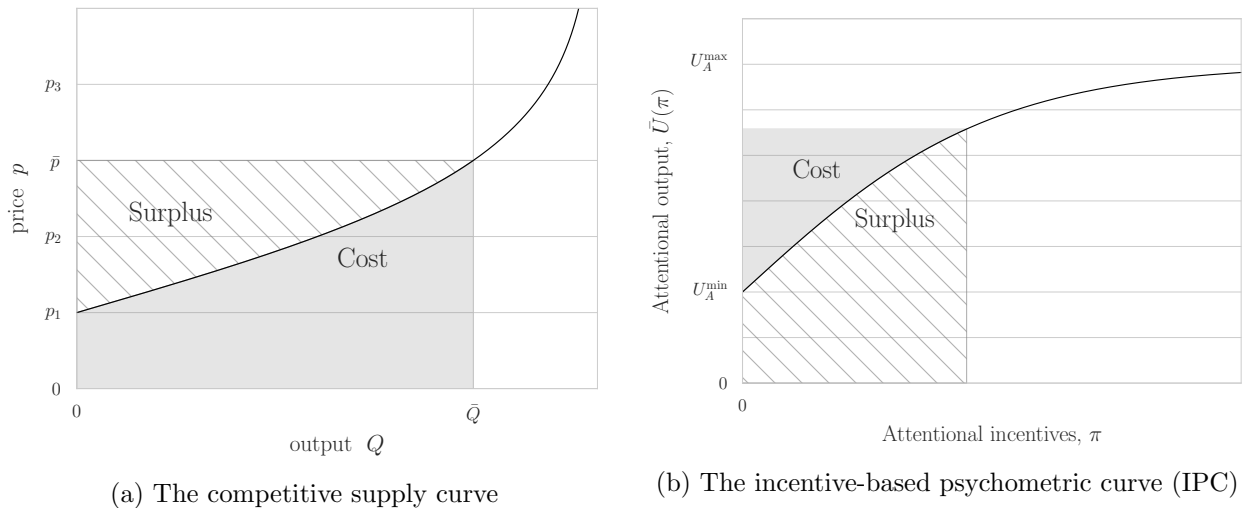


Figure 1: Analogy with competitive supply

We now specify the elements of the standard RI model of frictional mistakes in abstract terms. The decision maker (DM) is choosing from a finite set A of $N \geq 2$ actions. What they will receive as a result of each such action is uncertain, depending on some underlying state of the world, ω , which is ex ante unknown but in principle knowable from available information. The overall (finite) set of states is Ω . To know the state is to resolve all relevant uncertainty. Formally, action a yields utility $u(a, \omega) \in \mathbb{R}$ to the DM in state ω . The dependence of the utility function on the unknown state provides the incentive to investigate the available information and thereby become better informed.

RI theory models the DM as learning to a point that optimally balances the benefits of improved information against the subjective costs of becoming better informed. In technical terms, the first

important point is that the DM's subjective prior beliefs $\mu \in \Delta(\Omega)$ about the likelihood of the states interacts with the variability of utility across states to impact the possible benefits to learning. There is little incentive to learn if the knowledge is not likely to significantly update which action is best to take. In the example of medical insurance, prior beliefs concern the levels of service available under each plan and their costs. For example, if the DM believes that all insurance companies provide ready access to essentially identical medical services, they are unlikely to dedicate much attention to what they see as differences in fine details. Even a DM who believes that large differences in quality are likely, yet does not feel competent to judge these differences, may essentially ignore them and focus instead on costs. On the other hand if a DM is relatively aware of medical factors yet relatively innumerate, they may not be competent in comparing the costs of different plans. In this case they might use relatively crude prior beliefs about these costs in making choices. Hence different prior beliefs and mental frictions will give rise to systematically different patterns of mistakes.

By way of further illustration of how to apply the model in practice, consider the case of school applications. There has been much concern with the apparent fact that low income families appear to place lower weight on academic performance when choosing schools than do high income families [Hastings and Weinstein, 2008]. The open question has been the extent to which this reflects differences in priorities as opposed to differences in costs of gathering and processing information. If these costs are high, low income parents may “choose schools based on easier-to-determine characteristics such as proximity, instead of student test scores” [Hastings and Weinstein, 2008, p. 1374]. In a simple corresponding two attribute model, the prior is over how close schools are and how academically strong they are. Parents may find close-by schools trivial to identify, with academic quality of more distant schools harder to assess. In broad confirmation, Hastings and Weinstein [2008] show that direct presentation of relevant test information resulted in significantly more parents choosing higher-scoring schools for their children. Reducing frictional elements of learning materially impacts behavior, in line with the RI model.

As a third example consider the well-known phenomenon of failure of many of those who are eligible to claim the Employed Income Tax Credit [Currie, 2008]. One open question is the extent to which this is related to possible stigma as opposed to misinformation or uncertainty about eligibility and the claiming process. With regard to the latter, relevant prior beliefs concern eligibility and the extent of

the credit if eligible. To resolve this prior uncertainty requires some understanding of the tax code and of the claiming process. The evidence of [Bhargava and Manoli \[2015\]](#) suggests that this is not easy for all without guidance. They run interventions that reduce frictions associated with understanding the benefits to claiming and accurately completing the required forms. Their field study provides evidence that increasing the salience and the simplicity of information provision significantly increases the claiming rate. By way of contrast, stigma-reducing interventions appear to have no effect.

2.2 Attentional Incentives, π

With the model objects in place, we introduce the general attentional problem and corresponding linear incentives that mimic the role of output price in competitive supply. The linearity of expected utility theory plays a critical role in allowing us to identify a simple and general method of scaling attentional incentives. Given decision problem A , we construct a family of decision problems A_π , indexed by $\pi > 0$. We enumerate actions in A by $a(n)$ for $1 \leq n \leq N$ and correspondingly each action in A_π by $a_\pi(n)$. Utility for each action is proportionate to that for action $a(n)$ with factor of proportionality π . This defines the equivalent of price for a firm.

Definition 1. *The **linear family of decision problems** generated by A is a set of decision problems, $\{A_\pi\}_{\pi>0}$, with A_π having N actions $a_\pi(n)$ with state dependent utility for $1 \leq n \leq N$ satisfying,*

$$u(a_\pi(n), \omega) = \pi u(a(n), \omega) \geq 0. \tag{1}$$

To generate A_π for $\pi \in (0, 1]$ one can make use of lottery prizes. Specifically, define action $a_\pi(n)$ to be a lottery. With probability π it yields a prize precisely as does action $a(n)$ while with probability $(1 - \pi)$ it yields a prize that has zero utility. Note that our use of probabilistic rather than monetary rewards is not essential. Our approach allows one to recover costs in the appropriate expected utility units by applying a risk aversion correction when monetary rewards are scaled up and down.

2.3 Attentional Output, $\bar{U}(\pi)$

We consider an observer who sees choice in each decision problem repeatedly and records the joint probabilities over actions and states. As in the standard random utility model, in practice this may often derive from data on a large population choosing relatively few times from the same menu. As in [Caplin and Martin \[2015\]](#) and [Caplin and Dean \[2015\]](#), we call this state dependent stochastic choice (SDSC) data. Decisions are not always perfectly matched to the state, so that mistakes are potentially present in the data.

Definition 2. *An SDSC data set $P \in \mathcal{P}$ specifies probabilities $P(a, \omega) \geq 0$ over any state-action pair consistent with the prior,*

$$P \in \mathcal{P} \iff \sum_a P(a, \omega) = \mu(\omega).$$

The set of all chosen actions is denoted by $A(P) := \{a \in A \mid P(a) > 0\}$. Given any particular choice set, A , we write $\mathcal{P}(A) \subseteq \mathcal{P}$ as all data sets that are in principle consistent with this set of actions.

The ideal data comprises one observation of SDSC for each $\pi > 0$. We let $P_\pi(n, \omega) \geq 0$ specify the probability of choosing each possible action $a_\pi(n) \in A_\pi$ in each possible state. For each $\pi > 0$ we use this dataset to compute the corresponding revealed expected utility,

$$U(\pi) := \sum_\omega \sum_n u(a_\pi(n), \omega) P_\pi(n, \omega). \tag{2}$$

We would expect $U(\pi)$ to be increasing in π for two reasons. First, an increase in π raises the payoff to each action in each state proportionately. Second, $P_\pi(n, \omega)$ might place more weight on actions with high payoffs as the larger differences in $u(a_\pi(n), \omega)$ across actions raise the incentive to choose actions appropriate to each state. Our goal is to record the quality of the decision making, which requires us to eliminate the first effect. To that end we divide expected utility by π . This is equivalent to recording the expected utility that results from using the strategy for decision problem A_π with the rewards associated with A . It is the pure reflection of how the incentive impacts choice quality, removing the direct effect of the scaled incentive. It therefore provides a unidimensional measure of

attentional output.

$$\bar{U}(\pi) := \frac{U(\pi)}{\pi} = \sum_{\omega} \sum_n \frac{u(a_{\pi}(n), \omega)}{\pi} P_{\pi}(n, \omega) = \sum_{\omega} \sum_n u(a(n), \omega) P_{\pi}(n, \omega). \quad (3)$$

This is the attentional output depicted in the IPC graph of 1b. In what follows, we assume that $\bar{U}(\pi)$ is non-decreasing in the attentional incentive π just as in figure 1b. In proposition A.1 we show that this is a property of rational inattention (RI) models.

2.4 Bounds

We turn now to defining the bounds on the IPC as depicted in figure 1b. U_A^{\max} is the expected utility when there are no information costs. In this case, the agent makes the optimal choice in each state. With regard to the lower bound, U_A^{\min} is the utility associated with inattention. In this case, the agent chooses the action that is unconditionally best according to the prior. $\bar{U}(\pi)$ is always weakly greater than U_A^{\min} . In the observed data for $\pi \in (0, 1]$ the bounds may be tighter, ranging from minimum value $\bar{U}(0)$ to maximum value $\bar{U}(1)$,

$$\bar{U}(0) := \lim_{\pi \searrow 0} \bar{U}(\pi) \geq U_A^{\min} := \max_{a \in A} \sum_{\omega} \mu(\omega) u(a, \omega); \quad (4a)$$

$$\bar{U}(1) \leq U_A^{\max} := \sum_{\omega} \mu(\omega) \max_{a \in A} u(a, \omega). \quad (4b)$$

The lower bound will be above U_A^{\min} to the extent that attention is given in an essentially effortless or habitual manner. The upper bound will be below U_A^{\max} to the extent that it is infeasible for the decision maker to fully comprehend reality at any reward level.

3 RI and the Cost of Utility

It is intuitive that the IPC's shape reflects costs. The curvature of \bar{U} can be linked to the cost of moving from one expected utility level to another. For example, in figure 1b we see that \bar{U} is relatively steep close to U_A^{\min} which means that a relatively small increment in the incentive increases the attentional

output substantially. To put it differently, increasing attention is relatively inexpensive. On the other hand, the relative flatness of \bar{U} at larger values of π means that increasing attention is expensive.

In this section we make this intuition precise by showing how the simple theory that information is chosen to optimally balance the costs and the benefits of learning allows one to recover costs of learning from the IPC. Again, there is a complete analogy with how one recovers costs of output from the competitive firm's supply curve under the assumption of profit maximizing behavior.

It is the theory that behavior reflects optimal choices that allows one to recover costs of production from the supply curve. One maintained assumption is existence of a specific technology for the production of output that is independent of the price of output. Assuming that the firm is profit maximizing, this technology, together with the prices of inputs, determines the minimum cost of producing any level of output, Q . In the simple differentiable case, the profit maximizing choice of a particular level of output at a given price reveals the marginal cost of that output level. Integrating up marginal cost allows total cost to be recovered.

To pursue this analogy and derive the cost of attention from the IPC likewise involves a theory to be imposed on the data. The theory is precisely analogous to the theory of profit maximization: that choices are made to balance costs against benefits with a cost function that is not affected by payoffs. This is the general theory of RI. In specifying this, we need to think about the domain of choice. For the competitive firm it is easy: it is the output level (or vector of output levels when there are many goods). In RI theory, the choice domain is more intricate. Indeed it can be specified many equivalent ways: as choice of a signal structure; as choice of a distribution of posteriors; and as choice of a joint distribution over actions and states consistent with the prior. For our current purposes, this last approach, introduced by Sims [1998, 2003] and Matějka and McKay [2015], is most convenient. In this sense the feasible set of strategies is precisely the same as the set of possible observations of SDSC.¹

The restrictive aspect of RI theory derives from the assumption that the cost function is independent of the level of payoffs. This is much as in the economic theory of production, in which technology does not per se depend on the prices of inputs, at least in static environments. This assumption how-

¹Since the prior is held fixed, one can map the joint distributions to a collection of conditional distributions corresponding to general statistical experiments.

ever, does not exclude the possibility that attention costs can depend on the framing of the decision problem itself, as in [Tversky and Kahneman \[1981\]](#). In fact assessing the impact of changes in the framing of options on costs of attention is an important area of theoretical and empirical investigation (see [Caplin and Martin \[2019\]](#)).

A common feature of attention costs is that obtaining more information implies higher attention costs. Below we define the notion of Blackwell informativeness [[Blackwell, 1953](#)], a concept widely used in economics, on the state-dependent stochastic choice (SDSC) strategy space that we use.

Definition 3. For any $P \in \mathcal{P}$ define the unconditional probability of action, a , and the revealed posterior [[Caplin and Martin, 2015](#)] at any action, a ,

$$P(a) := \sum_{\omega} P(a, \omega) \quad \forall a \in A(P); \quad (5)$$

$$\gamma_P^a(\omega) := \frac{P(a, \omega)}{P(a)} \quad \forall (a, \omega) \in A(P) \times \Omega. \quad (6)$$

For two SDSC datasets, $P, Q \in \mathcal{P}$, consistent with the same prior, P is Blackwell more informative than Q , if and only if,

$$\sum_{a \in A(P)} \phi(\gamma_P^a) P(a) \geq \sum_{a \in A(Q)} \phi(\gamma_Q^a) Q(a), \quad (7)$$

for every convex continuous function $\phi: \Delta(\Omega) \rightarrow \mathbb{R}$.

Note, that expected utility theory implies that choice based on a Blackwell more informative experiment always yields weakly higher expected utility from an ex-ante perspective.

An attention cost function, denoted by K , specifies a utility cost for each possible attention strategy. We make the standard assumption that K is increasing in the Blackwell order. We assume also that the inattentive strategy has no cost and allow the possibility that some attention strategies can be prohibitively costly. Given the attention cost, an RI model assumes that the DM in decision problem A_π solves the following optimization problem,

$$\max_{P_\pi \in \mathcal{P}(A_\pi)} \sum_{\omega} \sum_n u(a_\pi(n), \omega) P_\pi(n, \omega) - K(P_\pi). \quad (8)$$

Assuming the solution exists we denote any maximizing SDSC strategy as \widehat{P}_π .

3.1 The Utility Cost Curve

When studying the optimizing decisions of a competitive firm with fixed input costs, economists have found it of benefit to separate the optimizing decision into stages. First, the firm is seen as identifying from its technology how most cheaply to produce any given level of output. This defines the cost curve. The firm then chooses an optimal level of output by equating marginal cost with marginal revenue.

Such a perspective is of value in RI theory as well. However, the output in the attentional problem is considerably more complex than a single quantity produced. The attentional output is the attention strategy defining the distribution of chosen actions for each state. Its value is defined by the payoff structure characterizing the decision problem. In order to exploit the simple marginal cost equals marginal revenue logic for the attentional problem we define the attention cost of achieving a certain level of expected utility.

Definition 4. Given $P \in \mathcal{P}$, $U(P)$ is defined as the achieved expected utility

$$U(P) := \sum_{\omega} \sum_{a \in A(P)} u(a, \omega) P(a, \omega). \quad (9)$$

We call the cost function for the production of expected utility the **utility cost curve**. It is the minimal cost of achieving any expected utility level $u \in [U_A^{\min}, U_A^{\max}]$ in decision problem A .

Definition 5. Given K , decision problem A , and $u \in [U_A^{\min}, U_A^{\max}]$, the **utility cost curve** $\bar{K}_A(u)$ identifies the lowest cost of achieving this level of expected utility,

$$\bar{K}_A(u) := \inf_{\{P \in \mathcal{P}(A) | U(P) \geq u\}} K(P). \quad (10)$$

Note that $\bar{K}_A(U_A^{\min}) = 0$, since the inattentive strategy achieves U_A^{\min} . It is intuitive that Blackwell monotonicity requires that $\bar{K}_A(u)$ be non-diminishing in u . We prove this in appendix [A.1](#).

The above construction is familiar in information theory. The utility cost curve is a version of the information rate distortion function, which has been thoroughly studied in information theory for the case of Shannon mutual information [[Cover and Thomas, 2006](#)]. In the current application, the utility

function plays the role of the distortion measure and the cost function is a general version of mutual information.

3.2 Marginal and Total Cost

The problem of the competitive firm is to choose its output level given the price of output,

$$Q(p) := \arg \max_{Q \geq 0} pQ - C(Q). \quad (11)$$

Elementary microeconomics shows how to recover the total cost of producing output $Q(p)$ by equating marginal cost to price,

$$C'(Q(p)) = p.$$

In the well-behaved case, this relationship—together with the assumption that no production is costless—allows us to recover the total cost. Given \bar{Q} ,

$$C(\bar{Q}) = \int_0^{\bar{Q}} p(Q) dQ = \bar{p}Q(\bar{p}) - \int_0^{\bar{p}} Q(p) dp,$$

where $p(Q)$ is the inverse supply curve.

We follow the same logic in rewriting the RI problem to recover the cost of attention. Using the utility cost curve we can write the DM's problem (8) for any $\pi > 0$ as,

$$\max_{u \in [U_A^{\min}, U_A^{\max}]} \pi u - \bar{K}_A(u). \quad (12)$$

The key observation is that if a strategy $P \in \mathcal{P}(A)$ yields utility u , then the corresponding strategy $P_\pi \in \mathcal{P}(A_\pi)$ yields utility πu and vice versa. This implies that $\bar{K}_A(u) \equiv \bar{K}_{A_\pi}(\pi u)$. Hence one can treat all decision problems A_π as deriving from a choice in the original decision problem A with correspondingly scaled incentives.

If \bar{K}_A is differentiable and convex, we can therefore characterize the solution to problem 12, $\hat{u}(\pi)$,

by a first-order condition,

$$\bar{K}'_A(\hat{u}(\pi)) = \pi.$$

How does $\hat{u}(\pi)$ relate to the IPC? Assuming that the data reflects optimal choice, the gross utility levels should coincide, $\pi\bar{U}(\pi) = \pi\hat{u}(\pi)$.² Hence, the IPC plots the maximizer of problem 12 parametrized by π ,

$$\bar{K}'_A(\bar{U}(\pi)) = \pi.$$

The recovery theorem shows this logic to be completely general.

3.3 Recovery

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

$$\bar{K}_A(\bar{U}(\pi)) = \pi\bar{U}(\pi) - \int_0^\pi \bar{U}(t) dt. \tag{13}$$




Figure 2: Cost recovery and the IPC

The recovery theorem states that the cost of attention for any problem A_π in the linear family is equal to the area between the IPC and the normalized utility level. For the general cost recovery result there is no need to assume that the utility cost curve is well-behaved. We use results in convex

²This implies that the IPC can also be viewed as the convex conjugate of the utility cost curve.

analysis [Rockafellar, 1971] to prove the theorem in appendix A.2.³ Figure 2 illustrates the recovery and the net achieved utility in decision problem A_{π_0} .

Our recovery theorem relates to earlier work in the literature. Caplin et al. [2017] present a recovery result that requires stronger conditions and richer data. de Oliveira et al. [2017] produce a recovery result based on a rich data set involving choice between choice sets. By constructing the linear family our approach provides a fully general recovery method that is also readily implementable.

Dewan and Neligh [2017] provide a recovery result analogous to ours for a restricted class of decision problems (“uniform guess tasks”) which are indexed by the utility of guessing the correct answer from a set of options. They show how to recover the cost of being correct from the relationship between the probability of being correct and the reward. In an experimental implementation they use their approach to test a variety of models, including those with Shannon costs and with fixed cost of attention. What allows us to produce a general recovery result that applies to all decision problems is the change of variables that we introduce in section 2 above.

There is also precedent in the psychological literature. In particular Musslick et al. [2018] established a recovery result that applies when DMs are risk neutral and the task is a guessing task à la Dewan and Neligh [2017]. This convergence of interest is part of a growing literature on costly mental labor incorporating the economic perspective to psychology [Botvinick and Kool, 2018].

3.4 Recovery from Finite Observations

Theorem 1 exploits rich variation in incentives for recovering attention costs. In most economic applications such fine variation in incentives is unavailable. Nonetheless, our approach still provides useful information about attention costs when behavior is only observed under a few incentive levels.

There are at least two potential paths forward. If we are willing to make further assumptions about the parametric class of the cost function we can estimate the most likely member of this class given the few observed points on the IPC. Using the estimate we can extrapolate the IPC to all incentive

³Note that the behavioral data generated by a utility cost curve and its convex hull are indistinguishable from one another and hence from data one can only recover the utility cost curve up to its convex hull. An analogous statement would apply if the cost function was not monotone in the Blackwell order. When data reflects optimal choice attention strategies observed under higher incentive levels are always weakly Blackwell more informative. Hence, the recovered utility cost curve is weakly increasing in the Blackwell order.

levels and use the integration in Theorem 1 to recover costs. This is the approach we take in our experimental analysis.

Without parametric assumptions on the shape of the cost function we have to give up obtaining a point estimate of the attention costs. However, we can derive bounds on it. The bounds are implied by the revealed optimality of the IPC—denoted by \bar{U} . It states that for any $\pi > 0$ we have,

$$\pi \bar{U}(\pi) - \bar{K}_A(\bar{U}(\pi)) \geq \pi v - \bar{K}_A(v) \quad \forall v \in [U_A^{\min}, U_A^{\max}].$$

Formally, the above relationship is tantamount to π being a subderivative of \bar{K} at $\bar{U}(\pi)$. We exploit these subderivative conditions to bound attention costs.

Specifically, assume that we observe behavior under M distinct incentive levels π_m with $\pi_{m+1} > \pi_m$ and $\pi_1 > 0$. To define the lower bound, which we denote as \bar{K}_A^{\min} , we use the tightest subderivative condition depending on the EU level in question.

$$\bar{K}_A^{\min}(u) := \begin{cases} 0 & \text{if } u < \bar{U}(\pi_1); \\ \pi_k (u - \bar{U}(\pi_k)) + \sum_{m=1}^{k-1} \pi_m (\bar{U}(\pi_{m+1}) - \bar{U}(\pi_m)) & \text{if } u \in [\bar{U}(\pi_k), \bar{U}(\pi_{k+1})]. \end{cases} \quad (14)$$

In order to derive the upper bound, which we denote as \bar{K}_A^{\max} , we use the sequential application of the subderivative condition together with the initial condition that guessing is costless, i.e. $\bar{K}_A(\bar{U}_A^{\min}) = 0$. For π_1 , using the point $v = \bar{U}_A^{\min}$ within the subderivative condition at $\bar{U}(\pi_1)$ we derive the upper bound,

$$\pi_1(\bar{U}(\pi_1) - U_A^{\min}) \geq \bar{K}_A(\bar{U}(\pi_1)).$$

Using this logic, we sequentially apply the subderivative condition together with the derived bounds to obtain the full upper bound.

$$\bar{K}_A^{\max}(u) := \begin{cases} \pi_k (u - \bar{U}(\pi_{k-1})) + \sum_{m=1}^{k-1} \pi_m (\bar{U}(\pi_m) - \bar{U}(\pi_{m-1})) & \text{if } u \in (\bar{U}(\pi_{k-1}), \bar{U}(\pi_k)]; \\ \infty & \text{if } u > \bar{U}(\pi_M). \end{cases} \quad (15)$$

The finer the grid of incentive levels under which we observe behavior, the tighter are the derived bounds on attention costs.

We implement our approach experimentally and discuss its applicability in various applied settings below.

4 The Microeconomic Toolbox

The fact that costs of attention play precisely the same role in consumer choice as do a competitive firm’s costs of production in its supply decision enables us to import the elementary microeconomic toolbox wholesale. In this section we show how to test for the existence of any rationalizing cost function. When such a function exists, we show how to recover consumer welfare net of costs precisely as one would estimate producer surplus from the supply curve. We also introduce a “revealed more costly than” binary relation that extracts rich information from each experimental observation. Finally, we show how to go back and forth between the IPC and the cost function.

4.1 Testing RI Theory

Just as in the case of revealed profitability conditions for the profit maximizing firm [Varian, 2014, Chapter 19], one can check in data whether the DM behaves in line with the model of rational inattention. There are two conditions that correspond to applicability of the RI model and rational expectations. One is “No Improving Action Switches” [Caplin and Martin, 2015], which insists on Bayesian expected utility maximization. The other condition, “No Improving Attention Cycles” [Caplin and Dean, 2015], rules out switching attention strategies across problems in a manner that increases overall utility. We restate them in a simpler form that also makes them independent.

Axiom A1. No Improving Action Switches: *Given A and corresponding P ,*

$$a \in A(P) \implies \sum_{\omega} P(a, \omega) u(a, \omega) = \max_{\tilde{a} \in A} \sum_{\omega} P(a, \omega) u(\tilde{a}, \omega).$$

Axiom A2. No Improving Attention Cycles: *Given a finite sequence of decision problems and*

corresponding SDSC data sets,

$$(A_m)_{1 \leq m \leq M},$$

with $A_1 = A_M$,

$$\sum_{m=1}^{M-1} \left(\sum_{a \in A_m} \max_{\tilde{a} \in A_m} \sum_{\omega} P_m(a, \omega) u(\tilde{a}, \omega) \right) \geq \sum_{m=1}^{M-1} \left(\sum_{a \in A_{m+1}} \max_{\tilde{a} \in A_m} \sum_{\omega} P_{m+1}(a, \omega) u(\tilde{a}, \omega) \right).$$

The two axioms are necessary and sufficient for the general RI model to apply. The No Improving Action Switches condition is a rationality condition for choosing actions given the available information. It states that for any action a that is selected with positive probability in the data, it must be that the corresponding expected utility is maximal under the revealed conditional distribution over states. There is no other available action which, if wholesale replaced action a , would yield a higher expected utility. The No Improving Attention Cycles condition, on the other hand, is a rationality condition for acquiring information. The fact that the set of available strategies is invariant to such features of the decision problem as the rewards to available options lies behind this axiom. Effectively it reduces to the statement that more of the ex ante uncertainty concerning utility is resolved the larger the incentive to do so. In this sense the attention strategies revealed under various decision problems have to reflect the associated incentives. Overall gross expected utility can not be improved by reshuffling attention strategies across decision problems. [Dean and Neligh \[2019\]](#) introduce statistical tests of these conditions for binary decision problems which we apply in our experimental data.

4.2 Welfare

When the data pass tests of RI theory, the IPC decomposes gross expected utility into the cost of information and the surplus. This provides a means for assessing the welfare implications of different paths to achieving a given decision quality. The same level of decision quality in terms of expected utility may be achieved through very different “technologies” of attention and information processing. Just as higher sales revenue does not imply higher profits, judging simply by the quality of the information reflected in the final decision may be misleading since it only takes into account the gross utility and ignores costs.

Figure 3 illustrates two IPCs which yield the same level of expected utility at π_0 , but with very different implications for net utility. Since $\bar{U}_1(\pi)$ lies above $\bar{U}_2(\pi)$ for all $\pi < \pi_0$, the marginal cost of information is greater over this range and the net welfare corresponding to \bar{U}_1 is greater than that corresponding to \bar{U}_2 .




Figure 3: Net welfare

4.3 More Informative and Revealed More Costly

Based on the analogy with elementary production theory and revealed profitability (see [Varian \[2014\]](#)) we introduce the notion of an unobserved strategy being “revealed more costly” than a chosen strategy. Given a rationalizing attention cost function, if an arbitrary strategy achieves a higher EU level than an observed—and hence optimally chosen—strategy in a given decision problem, then necessarily that strategy must have been more costly.

Definition 6 (Revealed more costly). *Given decision problem A , an arbitrary strategy $Q \in \mathcal{P}(A)$ is revealed more costly than an observed strategy $\hat{P}_A \in \mathcal{P}(A)$ if it achieves a weakly higher EU level.*

To illustrate the revealed more costly relation consider a symmetric binary decision problem with two equally likely states and payoff structure given by $u(a(1), \omega_1) = u(a(2), \omega_2) = 1$ and zero otherwise. The line running through point \hat{P}_A in figure 4 below denotes the iso-utility curve corresponding to the utility achieved by strategy \hat{P}_A . The slope of the iso-utility curve is -1 since the payoffs and the prior are completely symmetric. The gray region lying to the right of the iso-utility curve corresponding to \hat{P}_A collects all attention strategies that yield weakly higher expected utility. As a result, any other strategy interior to the gray region, for instance strategy \bar{Q} , must have been more costly.




Figure 4: Revealed more costly strategies

Note that the notion of revealed more costly is not vacuous—it is more restrictive than the Blackwell order which ranks all strategies to the north-east of \hat{P}_A as more informative as marked by the dashed line in the graph.⁴ One can go beyond the qualitative binary relation and provide quantitative bounds that apply to all unobserved strategies. The difference between the costs must be at least as big as the difference in the corresponding expected utility levels as depicted by the iso-utility curves. Hence, in figure 4 any rationalizing cost function must satisfy that for all Q lying in the shaded area,

$$K(Q) - K(\hat{P}_A) \geq U(Q) - U(\hat{P}_A).$$

This implies that using the incentive structure of the decision problem one can quantitatively compare costs of unobserved strategies with the costs of those that are used. In figure 4 the observed strategy yields expected utility level $U(\hat{P}_A) = .55$ while \bar{Q} yields $U(\bar{Q}) = .65$. For any rationalizing cost function, K , this implies a quantitative lower bound on the difference between the cost of these strategies,

$$K(\bar{Q}) - K(\hat{P}_A) \geq .1.$$

⁴In the binary case one experiment is more Blackwell informative than another if it lies to the north-east of it [Weber, 2010, p. 269].

4.4 From Theory to Data and Back: Example

Before moving to the experimental implementation we show how to move from theory to data and back from data to theory for the simple symmetric binary decision problem introduced above. The payoffs for decision problem A_π for $\pi > 0$ are,

	$a_\pi(1)$	$a_\pi(2)$	
ω_1	π	0	.
ω_2	0	π	

Suppose that the cost function is mutual information,

$$K(P) = \sum_{\omega} \sum_{a \in A(P)} P(a, \omega) \log \left(\frac{P(a, \omega)}{P(a)\mu(\omega)} \right).$$

Since the problem is entirely symmetric, the mutual information cost function implies that the optimal solution has to be symmetric as well. Correspondingly, to achieve any EU level u the SDSC strategy has to satisfy $P(a(1), \omega_1) = P(a(2), \omega_2) = u/2$.

Substituting in the cost function we get the following utility cost curve,

$$\bar{K}_A(u) = u \log(u) + (1 - u) \log(1 - u) - \log(0.5). \tag{16}$$

The implied IPC has the following logistic form,

$$\bar{U}(\pi) = \frac{e^\pi}{e^\pi + 1}.$$

Now consider the inverse problem in which we observe the logistic IPC. The recovery theorem

states that the cost of the attention strategy in problem π_0 can be recovered as,

$$\begin{aligned}\bar{K}_A(\bar{U}(\pi_0)) &= \pi_0 \bar{U}(\pi_0) - \int_0^{\pi_0} \bar{U}(t) dt. \\ &= \frac{e^{\pi_0}}{e^{\pi_0} + 1} \pi_0 - \int_0^{\pi_0} \frac{e^t}{e^t + 1} dt \\ &= \left(\frac{e^{\pi_0}}{e^{\pi_0} + 1} \right) \log \left(\frac{e^{\pi_0}}{e^{\pi_0} + 1} \right) + \left(\frac{1}{e^{\pi_0} + 1} \right) \log \left(\frac{1}{e^{\pi_0} + 1} \right) - \log(0.5).\end{aligned}$$

This expression coincides with the Shannon cost evaluated at the symmetric posteriors putting mass $\frac{e^{\pi_0}}{e^{\pi_0} + 1}$ on the more likely state. As the recovery theorem implies, this is precisely the utility cost curve associated with the Shannon cost in 16.

5 Experimental Design

In this section we illustrate our recovery result experimentally. The experiment relates to simple production tasks, such as identifying defective items on a manufacturing line, repeating a sewing pattern, or identifying bugs in software. [Guiteras and Jack \[2018\]](#) study a task of this form in field work in Zambia. Workers are paid for correctly classifying a set of beans. They run two treatments, monitored and unmonitored, and show that the number of errors of classification is significantly lower under monitoring.

Simple production tasks illustrate our approach in almost pure form. Since accurate completion requires attentional effort in the form of time and focus, mistakes are made. These can be and routinely are identified and may impact pay, as in the field experiment. The SDSC data set is therefore trivial to gather. The assumption that the prior is accurate seems reasonable for tasks that are routine and repeated thousands of times in any income generating setting. Moreover the linear incentive scheme can be implemented by probabilistic monitoring of performance, which is common in practice.⁵

⁵Interestingly, there are also marketing experiments that precisely match our linear family conditions by probabilistically implementing choice, asking whether individuals are willing to buy a product at a given price if there is an ex ante known probability π that they will get this opportunity (e.g. [Ding et al. \[2005\]](#), [Ding \[2007\]](#)). [Cao and Zhang \[2019\]](#) use a simple parametric model of attention costs to model the theoretical relationship between realization probability and elicited preferences. They use this to forecast demand in real purchase settings.

In our experimental production task we trivialize identification of the true state by using a psychometric design. Our experimental design implements the theoretical framework with a variety of incentives, while addressing a number of key practical considerations. After presenting the design and the experimental results, we return in section 7 to consider the value of our methods in field settings.

As illustrated in figure 5 our experimental task involves each subject being shown 24 geometric objects. Each shape is one of four regular polygons, seven- to ten-sided. The subject’s task is to determine whether there are more seven- or nine-sided polygons. The eight- and ten-sided polygons serve as decoys. Their presence lets us vary the total number of non-decoy shapes across rounds while still showing subjects 24 shapes in total. This ensures that counting only one of the non-decoy shapes is never sufficient to determine the realized state.

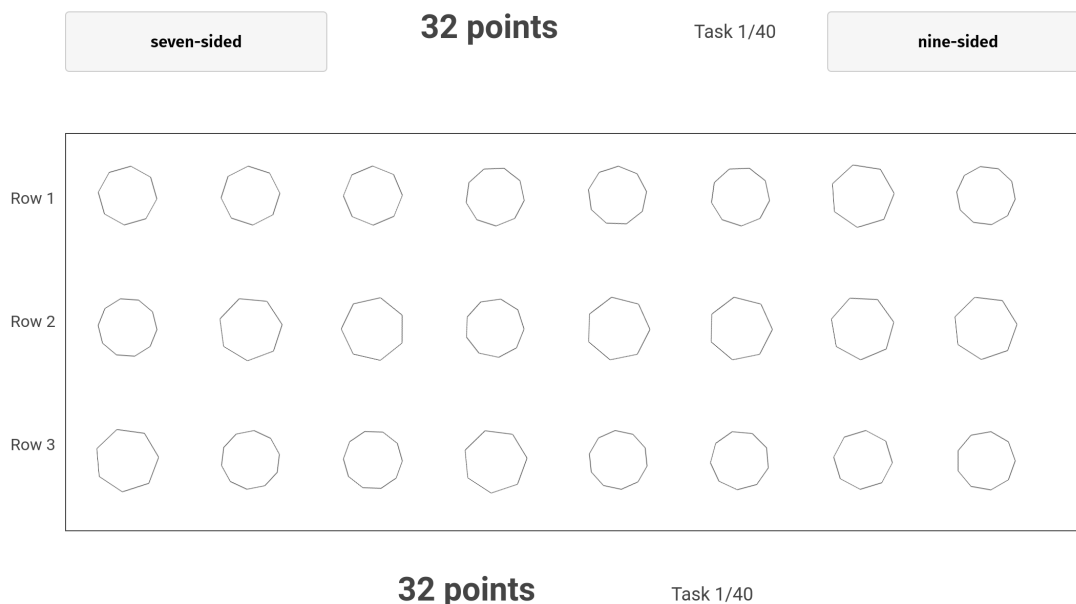


Figure 5: Example of an experimental task

The choice of the geometric counting task is motivated by its universality across cultures. Having both non-decoy polygons of the same parity is motivated by the symmetry in parallel sides—if one of the shapes has parallel sides while the other does not, differentiating between the two is relatively easy irrespective of the number of sides (as revealed in our pilots).

The location and rotation of the polygons as well as the total number of non-decoy shapes are generated randomly as described in appendix B.1. Hence, no round carries information about any

other round. In each round the non-decoy shapes are equally likely to be more common and subjects are told this in advance. Details of the general design can be found in appendix [B.2](#).

Our general design provides several channels that allow us to adjust the difficulty of the task. We can vary the total number of geometric shapes appearing on the screen; vary the difference between the numbers of shapes to be distinguished; and vary the number of sides of the polygons. The difficulty level of the task thus can range from trivial to virtually impossible.

In our experiment we varied the difficulty level of the tasks between subjects by setting the difference in the number of non-decoy polygons. The difference in the numbers of the non-decoy polygons was set at 1, 2, 3 and 6 constituting variation in our treatment of difficulty. Each subject faced tasks of a fixed difficulty level only.

The payment scheme is designed with the recovery theorem in mind. Subjects who make a correct decision are rewarded with points that count towards the probability of winning a fixed final prize of \$10. We vary the incentives according to a geometric scheme, the probability points can take values of 0, 1, 2, 4, 8, 16 and 32. The number of rounds for each incentive level is shown in the table below. Overall, the maximum score is 200.

Incentive level	Number of rounds
32 points	2
16 points	3
8 points	5
4 points	6
2 points	8
1 point	8
0 point	8
Total	40

Table 1: Incentive levels and number of rounds within a session

Subjects are not informed about their performance until the very end of the experiment, leaving them with probabilistic beliefs regarding the evolution of their score during the experiment. The effective incentives are thus set by the increments in the overall probability of winning. Since subjects only know the increments but not the base in their score, the design is robust to moderate non-linearities in expected utility as this form of averaging linearizes probability weights.

We wanted our subjects not to be rewarded for fully inattentive behavior. Since the polygons are generated randomly with either of the non-decoy shapes appearing in higher numbers with equal probability, the fully inattentive strategy of randomly answering in each round achieves a total score of 100 in expectation. To reward only attentive behavior the probability of winning the final prize of \$10 is determined by the formula $(\text{achieved final score} - 100)\%$. Subjects are informed of this.

Our design thus is a two factorial design varying both incentive and difficulty levels. Each subject generated data for one difficulty level and all of the incentive levels. We followed a between subject design and pooled across subjects when estimating attention strategies. Hence we recover the IPC of a representative agent. This gives us the necessary power to pin down parameters with sufficient confidence while leaving each subject to answer a moderate number of questions to avoid fatigue.

Our subject pool consists of U.S. workers on Amazon Mechanical Turk with at least 100 prior completed tasks with an overall approval rating of 95% or above. MTurk provides some desirable features for our experiment. First, subjects may have less intrinsic motivation to complete the tasks relative to an experimental lab. There is a clear opportunity cost of spending time on our experiment in terms of completing other tasks. The experiment is also scalable to a large sample size which helps us in providing the necessary statistical power for our analysis.

One challenge associated with the use of MTurk is our reliance on probabilistic rewards. Rewarding our subjects with lotteries through an on-line platform requires the implementation of a credible randomization device. In order to make the outcome of the lottery credible we let subjects stop the built-in clock of their computer and record the time up to millisecond precision. The final clock device is depicted below.




Figure 6: Example of the final clock device

The last two digits of the stopped clock—which are impossible to control intentionally—provide a credible uniform distribution which is then used to implement lotteries based on subjects’ final scores. For instance, if a subject achieves a final score of 167 points, then the subject has $(167 - 100)\% = 67\%$

chance of winning the final \$10 prize. This is implemented using the clock device. If the number based on the last two digits of the clock the subject is asked to stop upon completion is below 67 the subject wins the prize. Otherwise they win nothing. In the example shown in figure 6 the subject has to score 186 or above in order to win the final prize. Instructions and other screen-shots of the experimental interface are shown in appendix B.2.

We test the clock device in appendix C.3 and confirm that subjects had no control over stopping the clock and that empirical probabilities of winning the final prize generated by stopping the clock were aligned with the achieved final score.

5.1 Attention Cost Heterogeneity and Response Times

We use a between subject design to recover the IPC of a representative subject. Hence, we wish to confirm that the subjects we are pooling are homogeneous in their behavior to a satisfactory degree. Taking a look at the distribution of response times we see that a non-negligible fraction of subjects clicked through the whole experiment essentially participating in a low odds gamble. This pattern is indicative that the mental cost of an answer better than guessing is prohibitively costly to some of our subjects. On the other hand, many subjects perform better than guessing and respond to incentives. To accommodate the apparent heterogeneity in attention costs we split our subjects into two subsamples. One in which the representative behavior of subjects is not significantly better than guessing under any of the incentive levels and another in which the representative behavior is both better than guessing and responsive to incentives.

We find that total time spent on the experiment provides an informative indicator of attention costs. The representative behavior of subjects who spend less than 5 minutes in total on the experiment is indistinguishable from chance at each incentive level, while the behavior of those who spend more than that is both significantly better than chance and responsive to incentives as shown in figure 7. This threshold leaves us with 402 subjects out of the initial 814. In appendix C.1 we detail our analysis of response times and related performance. In appendix E we also conduct sensitivity analyses and we carry out our main analysis at 1-, 2-, and 12-minute thresholds in addition to the 5-minute one used in the main body of our paper. We find that our findings are qualitatively robust to the specific choice of threshold. The lower we set the threshold, the poorer the average performance, and hence

the higher the implied representative attention costs across the pooled subjects.

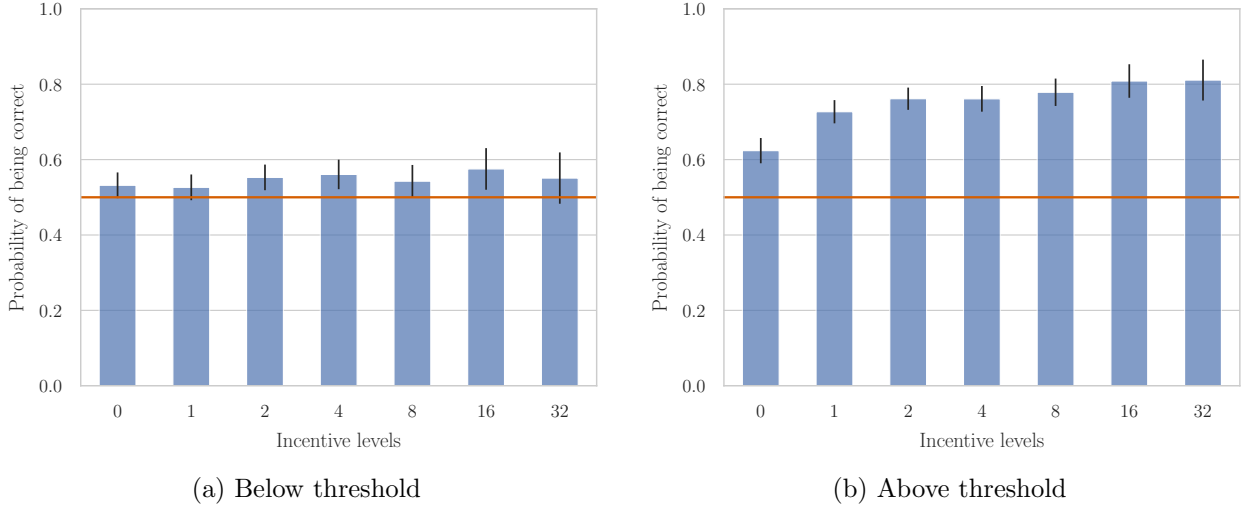


Figure 7: Average performance at different incentive levels

Note that performance is well above chance for the attentive sample even without any incentive.

5.2 Mapping the Experiment to the Theory

In this subsection we make explicit the mapping between the experiment and the theory. The state space is $\Omega = \{\omega_7, \omega_9\}$ depending on which non-decoy shape is more common. We set $\mu(\omega_7) = \mu(\omega_9) = 0.5$ and convey this to subjects. In each decision problem there are two actions, $A_\pi = \{a(7), a(9)\}$. The state dependent stochastic choices for a given incentive level, π , are denoted as $P_\pi(a(7), \omega_7)$ and symmetrically for nine-sided polygons.

For each task a subject can get reward points that count toward the probability of winning the final prize of \$10. In the experiment π can take values in $\pi \in \{0, .01, .02, .04, .08, .16, .32\}$. Hence, the utility function for a given incentive level is,

$$u_\pi(a(7), \omega_7) = u_\pi(a(9), \omega_9) = \pi u(\$10),$$

$$u_\pi(a(9), \omega_7) = u_\pi(a(7), \omega_9) = 0.$$

The IPC is therefore,

$$\begin{aligned}\bar{U}(\pi) &= \left(\mu(\omega_7)P_\pi(a(7) | \omega_7) + \mu(\omega_9)P_\pi(a(9) | \omega_9) \right) u(\$10) \\ &=: \bar{P}_\pi u(\$10).\end{aligned}\tag{17}$$

Equation 17 shows that in order to recover the IPC one only needs to estimate the average probability of being correct due to the symmetry of the payoff structure.

5.3 Experimental Test of RI

Given the symmetry of the design, for a fixed difficulty level, the No Improving Action Switches conditions collapse to,

$$P_\pi(a(7) | \omega_7) - P_\pi(a(7) | \omega_9) \geq 0;\tag{18}$$

$$P_\pi(a(9) | \omega_9) - P_\pi(a(9) | \omega_7) \geq 0.\tag{19}$$

Since the probabilities have to add up to one, the two conditions imply each other. Hence for each difficulty and incentive level, effectively, we have one condition to check. Four difficulty levels and seven incentive levels leave us with $4 \cdot 7 = 28$ instances to check the inequalities. When the incentive level is $\pi = 0$ the condition is vacuous, since without incentives rational subjects could pick arbitrarily. This means that effectively there are 24 non-vacuous conditions to check. We follow [Dean and Neligh \[2019\]](#) and compute the empirical counterparts of the inequalities to formulate statistical tests.

Considering the sample above the minimum time threshold, all 24 non-vacuous tests are passed at the 1% significance level. For the sample below the minimum time threshold only 12 out of the 24 tests are passed at the 1% significance level. We report p-values and t-statistics for both subsamples in appendix D.

Given the symmetry of the problem if the No Improving Action Switches conditions hold, the No Improving Attention Cycles condition implies that the average probabilities of being correct are increasing with the incentive level. RI theory doesn't put any restriction on how the different complexity

levels should compare to each other, hence we only test these conditions for one difficulty level at a time.

Note that in our setting the No Improving Action Switches condition requires only that on average the better action be picked given a realization of the state. With more states and corresponding best actions the condition is harder to satisfy. Similarly, given the binary setting, the No Improving Attention Cycles condition requires only that the average probability of being correct is increasing with the incentive level. With more states and non-symmetric payoffs this condition would be more demanding.

We report the increment in average probabilities for both subsamples below. A positive increment implies that the corresponding condition is satisfied.

$\Delta\pi$	Difference			
	1	2	3	6
1 - 0	-.026	-.045	.037	.007
2 - 1	.014	.041	.029	.024
4 - 2	.015	.022	-.026	.029
8 - 4	.019	-.041	-.002	-.062
16 - 8	-.016	.055	.029	.075
32 - 16	-.072	-.026	-.008	.018

Table 2: below 5-min. minimum time threshold

$\Delta\pi$	Difference			
	1	2	3	6
1 - 0	.108	.121	.056	.123
2 - 1	.050	.044	.042	.002
4 - 2	-.013	.000	.008	.005
8 - 4	.010	.022	.012	.026
16 - 8	.023	.012	.049	.038
32 - 16	-.003	-.005	.017	.003

Table 3: above 5-min. time threshold

We check the adjacent pairwise comparisons for each difficulty level separately. This leaves us with 24 tests to check. As expected we can see that there are fewer violations of the condition in the sample spending at least 5 minutes on the experiment. In this subsample none of the violations is significant at the 5% level while seven of the passed tests are significant at the 5% level. On the other hand, for subjects below the time threshold there are 10 violations altogether. We report the significance level of these tests in appendix D.

Altogether, we confirm that the representative behavior revealed by the subsample above the minimum time threshold of the dataset is in line with the basic assumptions of the RI model.

6 Experimental Results

In this section we empirically recover the IPC and use it in various applications. The key step is estimating the relationship between the incentive level and the probability of a correct answer.

6.1 Estimating the IPC

To start our analysis we consider only one level of task complexity, the simplest task in our analysis, when the difference between the numbers of non-decoy shapes is set to 6. Figure 8 depicts the relationship between the probability of correct discrimination and incentives. Taking a look at the




Figure 8: IPC for the simplest task

intercept we see that even absent incentives subjects can distinguish the states more than 70% of the time. This finding is in line with the psychological perspective in which the ability to discriminate is studied without any explicit incentives. As in many psychological tasks, we find that increasing the extrinsic incentives also increases the probability of correct discrimination. The recovery theorem implies that one can read off attention costs and net welfare from the above graph.⁶

We estimate the IPCs across all difficulty levels jointly. Given the geometric incentive scheme we fit a logit model with a log-transform of the incentive scheme. This parsimonious model lends itself

⁶As noted before, since the payoff structure is binary and symmetric it is sufficient to estimate the average probability of being correct across states. We check this asymmetry in appendix F and find that conditional on the difference between the numbers of non-decoy shapes being large enough, the state with more 9-sided shapes is slightly easier to identify than the one with more 7-sided ones. Since the IPC only depends on the average probability this does not change our estimation strategy.

to interpretability and easy calculations for costs and net welfare.⁷ Specifically, we estimate a linear model for the log-odds ratio of the average probability of a correct answer given the covariates of observation i ,

$$\log\left(\frac{\bar{P}_i}{1 - \bar{P}_i}\right) = \alpha + \sum_{\text{Diff} \in \{2,3,6\}} \alpha_{\text{Diff}} D_{\text{Diff},i} + \beta * \log(\pi_i + 1) + \sum_{\text{Diff} \in \{2,3,6\}} \beta_{\text{Diff}} D_{\text{Diff},i} \log(\pi_i + 1), \quad (20)$$

where $D_{\text{Diff},i}$ is a dummy variable indicating whether for observation i the difference between the numbers of non-decoy shapes is Diff. The baseline difference in the model is Diff = 1. Note that the log transform of the geometric incentive scheme is shifted to accommodate the non-incentivized case ($\pi_i = 0$) and that π_i is measured in percentages.

The estimated IPCs for the four difficulty levels—given by the difference in the numbers of non-decoy shapes to be distinguished—are depicted in figure 9. We plot 95% confidence bands for each of the estimated IPCs.




Figure 9: IPCs for all difficulty levels

Note that, for all levels of task complexity, subjects perform above chance even without external incentives. The difficulty of the task does have an impact on the non-incentivized performance with

⁷Given our between subject design a fully non-parametric approach would compute the fraction of subjects with correct response at each incentive level. We plot these fractions in figure 9 for each difficulty and incentive level.

the probability of correct discrimination decreasing in difficulty. Note also, that for each difficulty level higher incentive levels result in a higher probability of being correct with diminishing returns to incentives. The joint differences across difficulty levels—intercept and slope—are significant for each pair of tasks except between the difficulty levels corresponding to non-decoy differences of 2 and 3. The sensitivity analysis shows that the results are robust to setting different time thresholds for the sample included in the estimation. We provide regression tables showing the sensitivity analysis and joint significant tests in appendix E.

While the slopes are similar across incentive levels, the intercepts are very different. How well attentionally-demanding activities are performed in the absence of incentives is an important task-specific factor that our method uncovers.

6.2 Cost and Welfare

For each difficulty level we estimate the relationship between the incentive level and the average probability of getting a task correct to compute the IPC. From this we recover net welfare and costs corresponding to the same gross achieved expected utility under the four difficulty levels. We normalize the utility of the final \$10 prize to 1 and that of no prize to zero.

Figure 10 plots the estimated IPCs. For each of the IPCs we mark the point on the curve that achieves 20% of the utility of the final prize, $u(\$10)$, and shade the area representing this gross utility. We see that the higher the difficulty level of the task, the higher the incentive required to achieve the 20% level of gross utility. The higher incentives then require lower precision—probability of correct discrimination—to result in the same gross utility.

We use the recovery theorem to decompose total expected utility into welfare and cost components. The cost of achieving a given gross expected utility level is given by the part of the rectangle that lies above the corresponding IPC as discussed in section 3. We quantitatively gauge the welfare trade-off between task complexity and incentives. Taking 20% of $u(\$10)$ in gross utility, table 4 shows the corresponding incentive levels, achieved probabilities of being correct, costs, and net welfare. Cost and net welfare are expressed as percentages of $u(\$10)$.

The patterns are clear—the cost of achieving a level of gross utility is monotonically increasing in




Figure 10: Estimated IPCs with equal gross utility (20% of $u(\$10)$)

	Incentive π (%)	Probability correct	Cost in $u(\$10)$	Net Welfare in $u(\$10)$
1 difference	26.43	75.66%	1.09%	18.91%
2 difference	24.29	82.33%	1.08%	18.92%
3 difference	23.69	84.43%	0.92%	19.08%
6 difference	21.77	91.83%	0.77%	19.23%

Table 4: Net benefits of achieving 20% gross utility (in terms of $u(\$10)$)

the difficulty. Hence, net welfare is falling.

6.3 Psychometrics

In traditional psychometric experiments the probability of correct discrimination is plotted against stimulus levels instead of incentive levels. In fact, in the majority of traditional psychometric experiments subjects are not incentivized at all. Our experimental design allows us to plot the traditional “stimulus-based” psychometric curve varying the difference between the number of non-decoy shapes. Figure 11 plots this traditional psychometric curve using our experimental data.

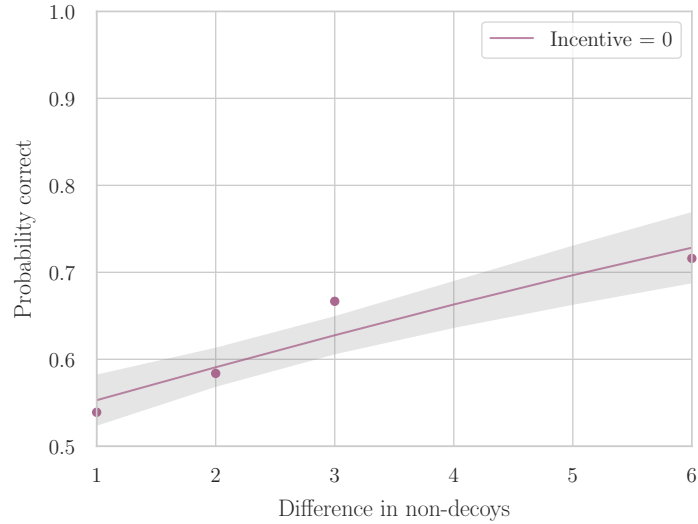


Figure 11: The stimulus-based psychometric curve with no incentive

Figure 12 plots the stimulus-based psychometric curves as a function of discriminability for each incentive level. The lowest curve corresponds to the traditional non-incentivized case. As expected we see that in all cases as differentiating between states gets easier the probability of giving a correct response gets higher. Furthermore, higher incentive levels shift the whole curve higher.

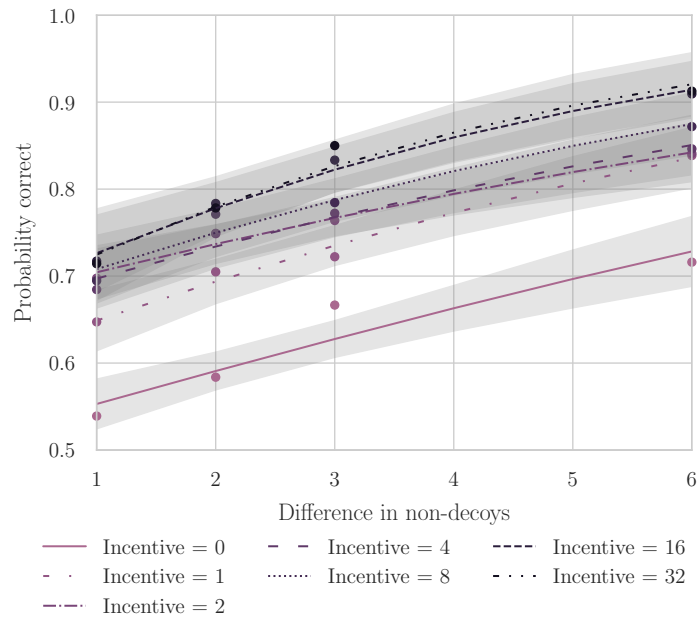


Figure 12: Stimulus-based psychometric curves under different incentive levels

Overall our experiment allows us to estimate the probability of a correct response as a function of

both the incentive and difficulty level.

Figure 13 plots the relationship between economic and psychophysical approaches to explaining the probability of correct discrimination. By varying both task complexity and incentives in our experiment we are able to quantify the trade-off using iso-performance curves, as depicted in figure 13b.

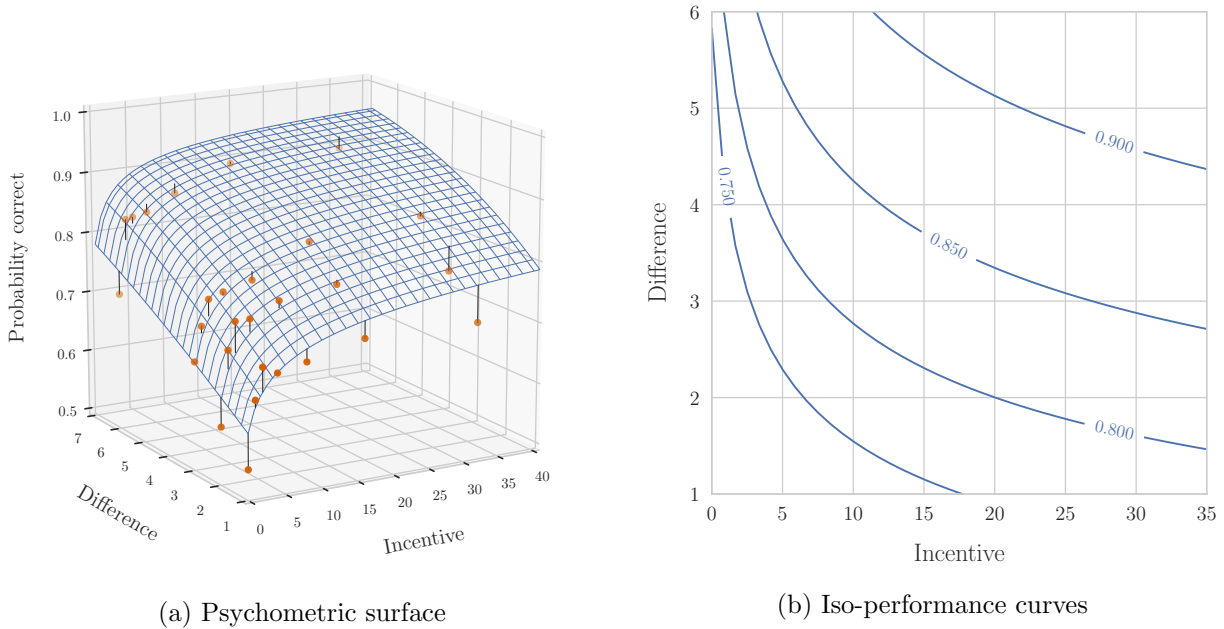


Figure 13: Trade-off between incentive and task complexity

As we see in the graph, if task complexity increases the DM has to be compensated with higher incentives to stay at the same performance level. In principle, analogous curves can be estimated for all psychometric tasks.

6.4 Attentional Inputs and Their Cost

Our recovery theorem quantifies costs of attention purely from observed behavior without specifying potential inputs entering the cost function. An important question is how our recovered cost relates to psychological inputs that can be thought of as sources of attentional cost—e.g. time and neural activity. In particular, time may be an important common source of attention costs in distinct attentional tasks as in the “mental labor theory” of Botvinick and Kool [2018]. This is analogous to labor being an important input factor in many production processes.

We correlate the recorded response times and our behavioral measure of cost to shed light on time as a costly input to attention. For robustness we take both the mean and median response times of subjects for each difficulty and incentive level and correlate them with the recovered cost using the IPC.⁸

Table 5 shows the correlations across all these variables.

	probability correct	mean time	median time	cost from IPC
probability correct	.	0.436	0.560	0.412
mean time	.	.	0.933	0.797
median time	.	.	.	0.880
cost from IPC

Table 5: Correlation between probability correct, time, and estimated cost

We see that the computed cost correlates highly with both measures of response time. As the table shows there is a significantly stronger correlation between our behavioral cost measure and time than between the probability of correct discrimination and time. Figure 14 plots the relationship between median and mean response time per round and the recovered costs for each incentive and difficulty level. It is intriguing to see that significant time is put into the experimental task even in the absence of incentives. This interplay between incentives and time suggests that the disutility (or utility) of the time input into attention may be very different between tasks and important to measure.




Figure 14: Correlation between estimated cost and time taken

⁸Since the probabilities of being correct in the data do not lie exactly on the IPC we use the predicted value from the IPC directly.

7 Applications

Our two key innovations are the rich data set (SDSC) and the scaling of incentives that allows cost recovery from this data. Applications that are essentially ideal for our approach are simple production tasks of which our experimental design is a bare-bones example. In these tasks, an observer can see both the actual choice and the correct choice as well as the payoffs to both. Our methods apply directly to these cases since SDSC is the natural data set and incentives can be linearly scaled by variation in the reward for successful task completion.

Production tasks are of clear economic importance and have been richly studied in the development context since they directly link attention with income. [Mullainathan and Shafir \[2013\]](#) have argued that poverty impacts income earning capacity through a variety of channels, including depletion of attentional resources. In support, [Kaur et al. \[2019\]](#) paid poor subjects a piece rate for completing a repetitive yet intricate attention-demanding task. They found that workers who were paid earlier were better able to complete these tasks and consequently earned higher income. They argue that their findings result from a direct relationship between financial constraints and worker productivity and that “psychological channels mediated through attention play a role in this relationship”. Subsequent research has identified a number of specific poverty-related factors that impact attentional constraints, such as sleep deprivation [[Bessone et al., 2019](#)], malnutrition [[Schofield, 2014](#)], exposure to high levels of air pollution [[Chang et al., 2016, 2019](#), [He et al., 2019](#)], and exposure to high levels of noise pollution [[Dean, 2019](#)].

This body of work raises the important issue of whether there may be a vicious cycle of poverty. While a low-skill/low-wage technology may currently be more profitable for producers, a more attentionally demanding technology might in principle pay for itself if the income-induced shift in human attentional capacity is large enough. This is a quantitative rather than a qualitative matter and requires improved measurement methods. Yet prior studies have done little to quantify the impact on attentional constraints of higher income, better sleep, lower ambient noise levels, or living in a less polluted environment. Even the definition of attentional effects is not uniform. Much like us, [Kaur et al. \[2019\]](#) measure attention by computing the probability of correct completion of a production task, while [Bessone et al. \[2019\]](#) and [Dean \[2019\]](#) use speed and accuracy on a stand-alone lab task as

their measure of attention.

Our methods apply directly to these cases. SDSC is the natural data set whenever the rules for classifying productive outputs as achieving or falling short of quality standards are clear cut. Moreover attentional incentives can be, and frequently are, scaled up and down either by variation in the reward for successful task completion or by variation in the probability of monitoring quality. Using our cost recovery method then allows measurement of attentional effects to be made uniform across domains. The rich quantification that our approach provides allows separate cost curves to be uncovered according not only to the attentional technology of income generation, but also according to cost shifters, such as noise and hunger.

Other possible applications are those in which complex images, documents, and histories have to be processed and classified for decision purposes. Examples include medical diagnoses, legal verdicts, and credit ratings. In this class, while some mistakes are observable and clear, in other cases opinions can differ. In the legal setting, [Shavell \[2009\]](#), [Cameron and Kornhauser \[2006\]](#), and [Stephenson \[2010\]](#) have argued that error reduction is a key element in the design of the legal systems, with appeals processes a key case in point. In the medical case, a 1999 Institute of Medicine Report finds that more than 40,000 people in the U.S. die due to medical errors [[Donaldson et al., 2000](#)]. Highlighting the role of attention, [Linder et al. \[2014\]](#) detected the impact of attention-influencing factors, such as meal breaks, on the willingness of doctors to prescribe antibiotics. In such cases, a key first stage in research is to invest in data analysis and essentially get as close as possible to producing ideal SDSC data. A key research goal is to develop expert systems for decision support, an area in which rational inattention theory is already beginning to make inroads [[Hoiles et al., 2019](#)].

8 Concluding Remarks

We introduce and implement a simple method of recovering attention costs from choice data. Our incentive-based psychometric curve allows costs of attention and consumer welfare to be recovered just as one can recover the costs of a competitive firm. This link to competitive supply allows us to apply the standard microeconomic toolbox to problems of attention. We implement our recovery method experimentally, outline ongoing applications in the field, and link our work to the broader literature.

Taken together, the forces that can be studied under the broad umbrella term of attention constitute a very significant new branch of social science: a novel hybrid of economics and psychology with a strong real-time interaction between models and applications. Going forward, it will be important to allow for interactions between frictional and mental gap based attentional phenomena. Generalized versions of our methods that recover the relative importance of different sources of mistaken decisions are needed.

Andrew Caplin

Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
andrew.caplin@nyu.edu

John Leahy

Department of Economics and
Gerald R. Ford School of Public Policy
University of Michigan
3308 Weill Hall
735 S. State St. #3308
Ann Arbor, MI 48109
and NBER
jvleahy@umich.edu

Dániel Csaba

QuantCo Inc.
670 Carolina Street
San Francisco, CA 94107
daniel.csaba@quantco.com

Oded Nov

Tandon School of Engineering
New York University
6 MetroTech Center, LC 401
Brooklyn, NY 11201
onov@nyu.edu

A Proofs

A.1 Blackwell Monotonicity and the Utility Cost Curve

Proposition 1. *Given K , $\bar{K}_A(u)$ is non-diminishing on $u \in [U_A^{\min}, U_A^{\max}]$.*

Proof of Proposition 1:

Suppose to the contrary that there exists $U_A^{\min} \leq u_1 < u_2 \leq U_A^{\max}$ such that $\bar{K}_A(u_1) > \bar{K}_A(u_2)$.

Take

$$P_2 \in \arg \min_{\{P \in \mathcal{P}(A) | U(P) = u_2\}} K(P).$$

By the linearity of U there exists $t \in [0, 1]$ such that

$$U(tP_2 + (1-t)P^I) = u_1, \tag{21}$$

with P^I being the inattentive strategy. By Blackwell monotonicity of K ,

$$\bar{K}_A(u_2) = K(P_2) \geq K(tP_2 + (1-t)P^I) \geq \bar{K}_A(u_1), \tag{22}$$

which contradicts $\bar{K}_A(u_1) > \bar{K}_A(u_2)$. □

A.2 Proof of Theorem 1

Lemma 1. *\bar{U} is non-decreasing in π .*

Proof of Lemma 1:

Take $\pi_2 > \pi_1 > 0$ and suppose $\bar{U}(\pi_2) < \bar{U}(\pi_1)$, so that

$$\bar{U}(\pi_2) \pi_2 < \bar{U}(\pi_1) \pi_2. \tag{23}$$

Since, \bar{K}_A is non-diminishing we have,

$$\bar{U}(\pi_2) \pi_2 - \bar{K}_A(\bar{U}(\pi_2)) < \bar{U}(\pi_1) \pi_2 - \bar{K}_A(\bar{U}(\pi_1)), \tag{24}$$

which contradicts the optimality of $\bar{U}(\pi_2)$ (and hence that of P_{π_2}). \square

Proof of Theorem 1:

For any $\pi > 0$, given state-dependent stochastic choice data, P_π , define,

$$N(\pi) := \sum_{\omega} \sum_n u(a_\pi(n), \omega) P_\pi(n, \omega) - K(P_\pi). \quad (25)$$

Consider an arbitrary $\pi_0 > 0$. This is associated with net utility $N(\pi_0)$ and normalized expected utility $\bar{U}(\pi_0)$. Consider now any other $\pi > 0$. Given optimality of data at π , P_π must do at least as well as applying P_{π_0} ;

$$\begin{aligned} N(\pi) &\geq \sum_{\omega} \sum_n u(a_\pi(n), \omega) P_{\pi_0}(n, \omega) - K(P_{\pi_0}) \\ &= \pi \frac{U(P_{\pi_0})}{\pi_0} - K(P_{\pi_0}) \\ &= \pi \bar{U}(\pi_0) - \pi_0 \bar{U}(\pi_0) + \pi_0 \bar{U}(\pi_0) - K(P_{\pi_0}) \\ &= (\pi - \pi_0) \bar{U}(\pi_0) + N(\pi_0); \end{aligned} \quad (26)$$

where the second line follows from the definitions of $u(a_\pi(n), \omega)$ and $u(a_1(n), \omega)$, the third line involves adding and subtracting $\pi_0 \bar{U}(\pi_0)$, and the final line applies the definitions of $N(\pi_0)$ and $\bar{U}(\pi_0)$. It follows immediately that N is convex and that $\bar{U}(\pi_0)$ is an element of the subdifferential of N at point π_0 ,

$$\bar{U}(\pi_0) \in \partial N(\pi_0).$$

[Rockafellar \[1971\]](#) Theorem 24.2. states that given a convex function N and a non-decreasing function \bar{U} satisfying $\bar{U}(\pi) \in \partial N(\pi)$ for all π ,

$$N(\pi) = \int_0^\pi \bar{U}(t) dt + C, \quad (27)$$

for some constant $C \in \mathbb{R}$. This is a generalized version of the fundamental theorem of calculus.

To pin down C , note that,

$$\lim_{\pi \searrow 0} N(\pi) = \lim_{\pi \searrow 0} \{U(P_\pi) - K(P_\pi)\} = - \lim_{\pi \searrow 0} K(P_\pi),$$

and

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

Hence, $C = - \lim_{\pi \searrow 0} K(P_\pi)$. But since the inattentive strategy is always available and costless in order to maximize welfare we need that $\lim_{\pi \searrow 0} K(P_\pi) = 0$. We conclude that $C = 0$.

The definition of the IPC together with equation 27 and C being zero gives

$$\pi \bar{U}(\pi) - \int_0^\pi \bar{U}(t) dt = \pi \bar{U}(\pi) - N(\pi). \quad (28)$$

Finally, since expected utility must be divided between net utility and information costs we have,

$$N(\pi) = U(\pi) - K(P_\pi) = \pi \bar{U}(\pi) - \bar{K}_A(\bar{U}(\pi)). \quad (29)$$

Rearranging and using equation 3 gives

$$K(P_\pi) = U(P_\pi) - N(\pi) = \pi \bar{U}(\pi) - N(\pi), \quad (30)$$

$$\bar{K}_A(\bar{U}(\pi)) = \pi \bar{U}(\pi) - \int_0^\pi \bar{U}(t) dt. \quad (31)$$

which is the area above the IPC. □

B Experimental Design

The experiment is available at <http://distinguishshapes.herokuapp.com/welcome>.

B.1 Design

An experimental session consists of 40 rounds of randomly generated tasks of the same difficulty. A task of fixed difficulty is described by the following algorithm and hyper-parameters.

- N : the total number of polygons generated in each round
- n : the number of different polygons to be present on the page
 - specify their types – e.g. [7-gon, 8-gon, 9-gon, 10-gon]
- m : the number of non-decoy polygons which the subject actually has to consider
 - specify their types – e.g. [8-gon, 10-gon]
- $\Delta^{\max} \leq \frac{N}{n}(n - m)$: the difference between the highest number and any other shape within the non-decoy class. This difference is assigned randomly to one of the polygons in the non-decoy class.
- $\Delta^{\text{mud}} \leq \frac{1}{m} \left(\frac{N}{n}(n - m) - \Delta^{\max} \right)$: shifts all the counts uniformly in the decoy class with a number in the range $[-\Delta^{\text{mud}}, \Delta^{\text{mud}}]$ picked randomly. Call the realization of this random number δ^{mud} .
- With these specified parameters (and the random mudding) make the $(n - m)$ decoy polygon(s) so that everything sums to N and the decoys are “as even as possible”. That is look for the closest integer assignment to $\frac{N}{n} - \frac{\Delta^{\max} + m\delta^{\text{mud}}}{n - m}$.
- Generate random degrees for rotating the polygons.

Example:

- $N = 24$
- $n = 4$ [7-gon, 8-gon, 9-gon, 10-gon]
- $m = 2$ [7-gon, 9-gon] (non-decoys)

$$\left[\overbrace{6}^{\text{8-gon}}, \underbrace{6, 6}_{\text{7-gon, 9-gon}}, \overbrace{6}^{\text{10-gon}} \right]$$

- $\Delta^{\max} = 2 \leq \frac{N}{n}(n - m) = 12$. Then randomly assign $\Delta^{\max} = 2$ to one of the non-decoy polygons.

$$\left[\overbrace{6}^{\text{8-gon}}, \underbrace{8, 6}_{\text{7-gon, 9-gon}}, \overbrace{6}^{\text{10-gon}} \right]$$

- $\Delta^{\text{mud}} = 3 \leq \frac{1}{m} \left(\frac{N}{n}(n - m) - \Delta^{\max} \right) = 5$. Generate a random integer between $[-3, 3]$. Suppose the realization is $\delta^{\text{mud}} = 2$. Shift all the non-decoys.

$$\left[\overbrace{6}^{\text{8-gon}}, \underbrace{10, 8}_{\text{7-gon, 9-gon}}, \overbrace{6}^{\text{10-gon}} \right]$$

Now change the decoys, so that they are closest to $\frac{N}{n} - \frac{\Delta^{\max} + m\delta^{\text{mud}}}{n - m} = 3$. In the example, that implies $[3, 3]$. Now the final numbers are


$$\left[\overbrace{3}^{\text{8-gon}}, \underbrace{10, 8}_{\text{7-gon, 9-gon}}, \overbrace{3}^{\text{10-gon}} \right]$$

They sum to 24 as per algorithm.

B.2 Instructions

You are about to complete an experiment and potentially earn money based on your performance. By completing the experiment you will receive a participation fee of 5 cents; on top of the participation fee you can win a \$10 bonus based on your performance. Below is a description of the task you will be facing.

You will be shown a screen with 24 geometric shapes. There will be 4 different types of shapes, and these are: seven-sided, eight-sided, nine-sided and ten-sided. Below are the 4 geometric shapes you will encounter.




The Task: Even though there are 4 different types of shapes your Task relates to only two of them, seven-sided and nine-sided shapes. There will always be [*difference in non-decoys depending on the complexity of the session*] more than the other and your Task is to decide whether seven- or nine-sided shapes are present in larger numbers. In each of the Tasks you face, these two shapes are equally likely to be present in larger numbers.

Note that the eight-sided and ten-sided shapes on each screen are irrelevant to your Task. While all Tasks are designed to be roughly equally difficult, the total number of these irrelevant shapes will vary in a narrow range.

You will only receive feedback on whether you were correct or not at the very end of the experiment.

Example of the Task

Example. This example has only 5 shapes; in the real Task you will face a similar task with 24 shapes.



Please decide whether there are more seven-sided or nine-sided shapes on the screen.

seven-sided

nine-sided

You will face 40 of these Tasks during the course of the experiment. For each Task shapes are randomly generated by a computer, and placed randomly on an eight-by-three grid. As a result, no Task carries information about any other Task. You can spend as much or as little time on each Task as you wish, but you have to complete all 40 Tasks in order to receive the participation fee and potentially earn a bonus.

Reward

By completing the experiment you will receive a participation fee of 5 cents. You may also win a bonus of \$10 depending on the number of points you score. If you answer all Tasks correctly you will get 200 points, and you will have a 100% chance of winning the bonus of \$10. More generally, if you achieve more than 100 points, your probability of getting the bonus is given by the following formula: $(\text{total achieved points} - 100)\%$ —so if you achieve 167 points, your final score will be 67, and your probability of winning the bonus will be 67%. If you answer only one out of every two correctly, as you could by guessing, you will get 100 points. If you score 100 points or below, there is no chance (0%) that you will earn the bonus.

How to score points

In fact there are 7 different point levels that you can get for answering correctly: 32 points, 16 points, 8 points, 4 points, 2 points, 1 point and 0 points. In all cases, a wrong answer gives you 0 points, so that 32 point rounds involve the highest reward for answering correctly, while 0 point rounds the lowest reward for answering correctly.

The point level is known to you at the start of each Task: it will be bold in the center of the top line of the screen displaying the Task.

In total, there are 2 Tasks offering 32 points for a correct answer, 3 Tasks offering 16 points, 5 Tasks offering 8 points, 6 Tasks offering 4 points, 8 Tasks offering 2 points, 8 Tasks offering 1 point, 8 Tasks offering 0 points. This means that in total, you will face 40 Tasks.

Note that the point scoring system and the experimental design explain why the maximum number of points that you can earn is 200 and why guessing would be expected to get you 100.

Winning the Bonus

To determine if you win the bonus, we will let you stop the clock below. This is your computer clock, presenting the time down to the millisecond (1/1000th of a second).

If the last two digits of the stopped clock are strictly less than your final score (this is your total achieved points - 100) you win the \$10 bonus, and if they are more, then you win nothing

Now, try to stop the clock showing the current time to millisecond precision. You should know that it is impossible for you to control these last two digits of the millisecond clock because of the time it takes the human brain and hand to respond. The purpose of this is to generate a random number, and match your probability of winning to your achieved score: the higher your score is, the more likely you are to win the \$10 bonus. To confirm this randomness, you can start and stop the clock as many times as you want, to check if you can stop it at a number of your choice. Please stop the clock at least 3 times.

17 :14:01:852

Stop the clock	Start the clock
Continue	

Your final score is the sum of all your points minus 100. The last two digits of the clock are 48. You win the prize if your final score is above 48.


Now, let's practice the Task. You will face 4 practice Tasks and then the final clock device to ensure familiarity. Remember, your Task is always to decide whether there are more of the seven- or nine-sided shapes. The shapes are surrounded by a black border. Please make sure that you can see this entire border before trying to complete the Task. Depending on your device you might have to scroll to see all of them.

seven-sided

32 points

Task 1/40

nine-sided



32 points Task 1/40

C Exploratory Data Analysis

Our design varies the difficulty level and the incentive level. The difficulty is set by the difference in the total number of non-decoy shapes. Table 6 shows the number of subjects for each difficulty level.⁹

Difficulty level (in difference of non-decoys)	Number of subjects
1	223
2	195
3	208
6	188
Total	814

Table 6: Number of subjects

For each subject we have 40 observations, hence, the total number of observations—an observation being a response to a task with certain difficulty and incentive level—is 32,560.

⁹Subjects who started but did not finish the experiment were dropped. There were 24 of them.

C.1 Decision Time Patterns

Even though our theory is silent on the relationship between time spent on a task and the corresponding precision level achieved it is revealing to explore time patterns in the data.

First, we plot the distribution of time spent on all 40 rounds in the experiment. Since MTurk requires a time limit on tasks, the limit was set at 3 hours, which provides participants ample time to go through every round slowly and carefully. No participant was timed out as a result of this restriction. Figure 15 demonstrates the histogram—the number of subjects spending more than 80 minutes on the experiment was 10 out of 814.




Figure 15: Distribution of time spent on the experiment

There is considerable variation in the total time spent on the 40 rounds of the experiment across subjects. This variation persists if we condition on difficulty level. In fact, many of the subjects seem to try their luck and click through the experiment.

95 subjects spent less than 1 minute on the whole experiment—all 40 rounds without the time spent on the instructions. The histogram of their final scores together with the corresponding average and 95% confidence interval is shown below. We can confirm that the inattentive strategy yields points matching the performance of flipping a coin as shown in figure 16.

As we increase the time threshold the average final scores start to increase as expected—our main trade-off in setting the threshold is separating the behaviorally distinct subjects with prohibitively high opportunity costs from the ones gradually adjusting their behavior while retaining power. We

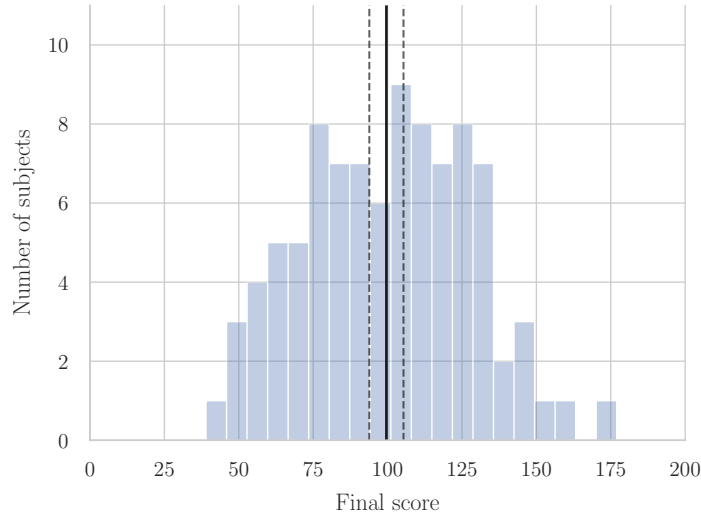
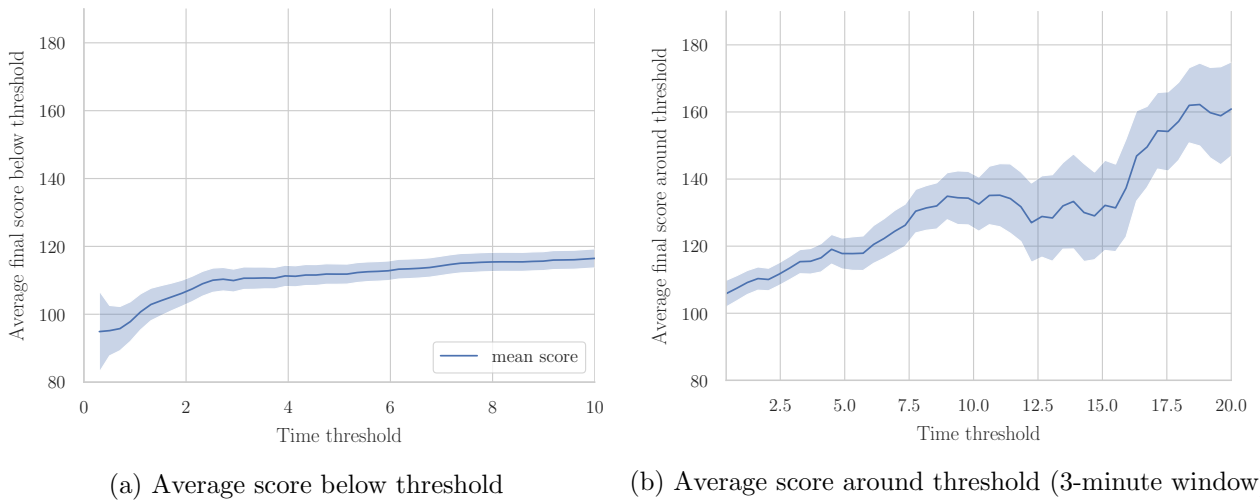


Figure 16: Histogram of final scores (<1 minute of total time)

plot the evolution of scores with time spent on the experiment in two ways. We graph the average final scores together with confidence intervals as we change the threshold—first, we take the overall averages below a given threshold, then we compute the averages for a 3 minute moving window. We see that below one minute of total time spent on the experiment the average score is not significantly different from zero. Raising the threshold increases the achieved score as expected.



(a) Average score below threshold

(b) Average score around threshold (3-minute window)

Figure 17: Expected final score and time spent on experiment

We set the time threshold to 5 minutes and carry out sensitivity analysis for thresholds varying between 1 and 10 minutes. The choice of threshold for splitting the sample according to total time spent on the experiment affects the results in a relatively stable and monotone way. The higher the

threshold the higher the probability of being correct for each decision level. We provide details of the sensitivity analysis with respect to the choice of the threshold in appendix E.1 when we present the estimated regressions.

The number of subjects who spend at least 5 minutes in total on the 40 rounds is 402.

Difficulty level (in difference of non-decoys)	Number of subjects
1	112
2	97
3	90
6	103
Total	402

Table 7: Number of subjects spending at least 5 minutes on the experiment

The representative behavior of subjects who spend less than 5 minutes on the experiment overall is not significantly better than pure guessing for any of the incentive levels.




Figure 18: Average probability of being correct in a round (spending < 5 minutes in total)

The non-responsiveness to incentives can be seen in the median time spent on each round conditional on the incentive level for subjects below the minimum time threshold. On the other hand, subjects above the threshold show considerable variation as shown in figure 19.




Figure 19: Median times per round for each incentive level

C.2 Learning and Fatigue

In order to test learning effects and fatigue, we study the dependence of average performance—the probability of being correct in a given round—on the round rank for subsamples both below and above the 5-minute minimum time threshold.

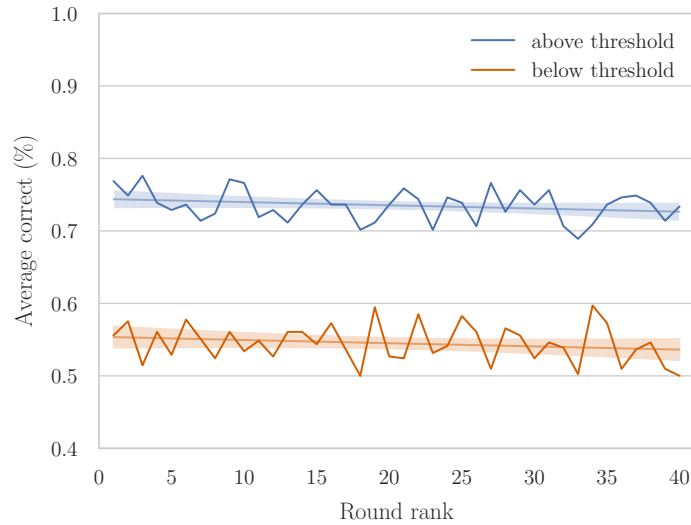


Figure 20: Average precision across rounds

As seen in figure 20 we find that there is no significant variation in performance across time for either subsamples. As expected the average performance levels are significantly different across subjects above and below the minimum time threshold. Gaining power from pooling across all difficulty

levels does make the average performance below the threshold significantly better than chance.

Table 8: Performance across rounds

	Above threshold	Below threshold
Intercept	0.744*** (0.010)	0.554*** (0.008)
Round	-0.000 (0.000)	-0.000 (0.000)
Adj. R-squared	0.00	0.00
No. observations	16080	16480

Clustered standard errors in parentheses.

* $p < .1$, ** $p < .05$, *** $p < .01$

Subjects above the minimum time threshold show a clear pattern of spending less time per round in later rounds. For subjects below the threshold median time plateaus at 1 second per round after the second round of the experiment.

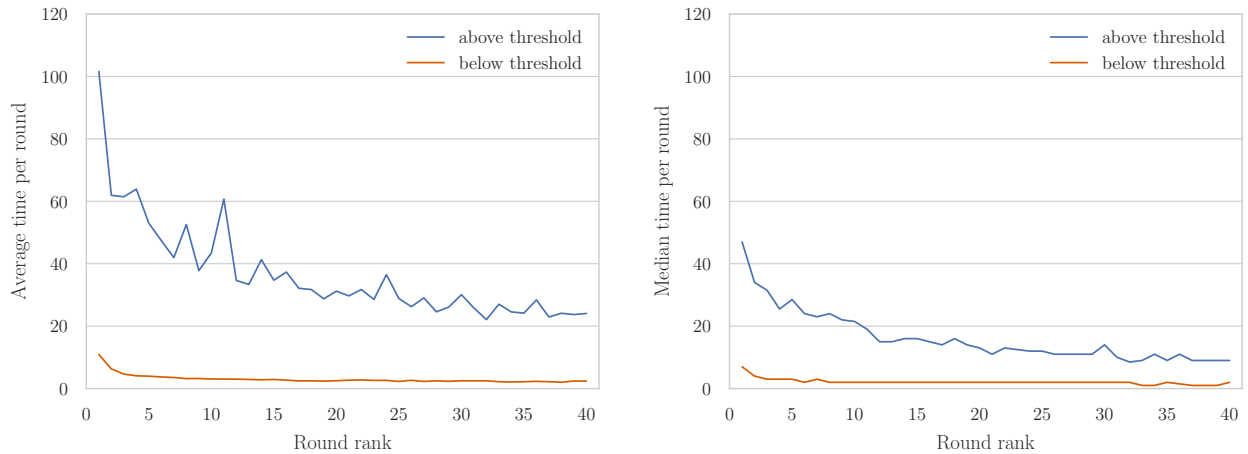


Figure 21: Average and median time across round ranks

Together with the roughly constant precision this is indicative that subjects might be getting better at the task as the rounds progress. However, they leverage this improvement by spending less time while keeping precision roughly at the same level.

C.3 The Clock Device and the Probability of Winning

To test the clock device conditional on being above 100 we check the expected number of winners just by looking at the final scores. Denoting the final score of subject i by x_i we have,

$$\text{expected number of winners} = \sum_{i|x_i > 100} \frac{x_i - 100}{100} = 320.86.$$

The corresponding standard deviation of the sum of independent Bernoulli random variables is 9.33. The actual number of winners in the sample is 313 which is within one standard deviation from the expectation and suggests that stopping the clock was effective as a randomization device.

D Testing RI Theory

We report p-values and t-statistics (in brackets) of the NIAS tests for both the attentive and inattentive subsamples below.

Difference				
$\Delta\pi$	1	2	3	6
0	.018 (2.09)	.012 (2.25)	.448 (0.13)	.001 (3.08)
1	.311 (0.49)	.598 (-0.25)	.010 (2.34)	.000 (3.66)
2	.089 (1.35)	.024 (1.99)	.000 (4.06)	.000 (4.82)
4	.028 (1.92)	.003 (2.82)	.015 (2.18)	.000 (5.51)
8	.003 (2.75)	.245 (0.69)	.030 (1.89)	.007 (2.46)
16	.069 (1.48)	.007 (2.45)	.005 (2.59)	.000 (4.32)
32	.837 (-0.98)	.095 (1.32)	.033 (1.85)	.000 (4.15)

Table 9: NIAS test p-values
Below minimum time threshold

Difference				
$\Delta\pi$	1	2	3	6
0	.008 (2.39)	.000 (4.59)	.000 (8.91)	.000 (12.40)
1	.000 (8.82)	.000 (11.41)	.000 (11.98)	.000 (19.55)
2	.000 (11.83)	.000 (13.83)	.000 (14.15)	.000 (19.58)
4	.000 (9.59)	.000 (12.03)	.000 (12.67)	.000 (17.25)
8	.000 (9.28)	.000 (11.93)	.000 (12.06)	.000 (16.88)
16	.000 (7.91)	.000 (9.68)	.000 (10.99)	.000 (14.39)
32	.000 (6.42)	.000 (7.78)	.000 (9.39)	.000 (11.85)

Table 10: NIAS test p-values
Above minimum time threshold

E Regression Results

E.1 Regression Results for the IPC

Table 11 presents logistic regressions of being correct as a function of the differences in the number of non-decoy shapes and the incentive levels. As in the figures the incentives are transformed as $\log(\pi + 1)$.

Table 11: Regression results for the IPC

	1 minute	2 minutes	5 minutes	12 minutes
Intercept	0.259*** (0.046)	0.299*** (0.052)	0.357*** (0.068)	0.493*** (0.104)
C(Diff.)[2]	0.099 (0.072)	0.121 (0.082)	0.189* (0.106)	0.194 (0.148)
C(Diff.)[3]	0.164** (0.075)	0.248*** (0.085)	0.387*** (0.118)	0.478*** (0.174)
C(Diff.)[6]	0.549*** (0.088)	0.623*** (0.100)	0.740*** (0.126)	0.841*** (0.177)
Incentives	0.134*** (0.029)	0.158*** (0.034)	0.235*** (0.045)	0.469*** (0.079)
C(Diff.)[2]*Incentives	0.030 (0.043)	0.043 (0.049)	0.073 (0.069)	0.001 (0.114)
C(Diff.)[3]*Incentives	0.018 (0.041)	0.015 (0.049)	0.061 (0.068)	-0.040 (0.118)
C(Diff.)[6]*Incentives	0.104** (0.049)	0.162*** (0.060)	0.190** (0.081)	0.161 (0.143)
Pseudo R-squared	0.01	0.02	0.03	0.05
No. observations	28760	23480	16080	10360

Clustered standard errors in parentheses.
 * $p < .1$, ** $p < .05$, *** $p < .01$

The results are stable and move monotonically with the time threshold across difficulty levels.

For the 5-minute threshold used in the main body of the paper the F-tests for testing the differences between the IPCs are shown in the following table.

	1-Diff.	2-Diff.	3-Diff.	6-Diff.
1-Diff.		.0372	.0004	.0000
2-Diff.			.2784	.0000
3-Diff.				.0026

Table 12: F-tests (p -values) for the difference between difficulty levels (5-min. threshold)

By way of comparison, we carry out the main estimation of the IPCs for the subject pool who spend less than 5 minutes overall on the experiment.




Figure 22: IPCs for the inattentive sample

Only the least difficult task—with the difference in the numbers of distinguishable shapes being 6—is significantly different from other difficulty levels.

	1-Diff.	2-Diff.	3-Diff.	6-Diff.
1-Diff.		.6862	.1899	.0001
2-Diff.			.6749	.0022
3-Diff.				.0085

Table 13: F-tests (p -values) for the difference between difficulty levels (below 5-min. threshold)

E.2 Regression Results for the Stimulus-based Psychometric Curve

Table 14: Regression results for the Stimulus-based Psychometric Curve

	1 minute	2 minutes	5 minutes	12 minutes
Intercept	0.046 (0.058)	0.053 (0.064)	0.058 (0.079)	0.166 (0.109)
C(Incentives)[1]	0.121 (0.077)	0.200** (0.086)	0.353*** (0.107)	0.507*** (0.155)
C(Incentives)[2]	0.362*** (0.076)	0.404*** (0.085)	0.649*** (0.104)	0.884*** (0.154)
C(Incentives)[4]	0.327*** (0.083)	0.385*** (0.093)	0.594*** (0.118)	0.973*** (0.174)
C(Incentives)[8]	0.336*** (0.092)	0.377*** (0.104)	0.616*** (0.137)	1.031*** (0.208)
C(Incentives)[16]	0.366*** (0.109)	0.440*** (0.125)	0.640*** (0.163)	1.252*** (0.241)
C(Incentives)[32]	0.294** (0.126)	0.313** (0.143)	0.613*** (0.193)	1.168*** (0.293)
Diff.	0.106*** (0.019)	0.126*** (0.021)	0.155*** (0.027)	0.164*** (0.037)
C(Incentives)[1]*Diff.	0.044* (0.023)	0.043* (0.026)	0.049 (0.034)	0.092* (0.049)
C(Incentives)[2]*Diff.	0.011 (0.023)	0.018 (0.027)	0.007 (0.033)	0.018 (0.048)
C(Incentives)[4]*Diff.	0.028 (0.024)	0.037 (0.028)	0.027 (0.036)	0.028 (0.055)
C(Incentives)[8]*Diff.	0.020 (0.028)	0.049 (0.032)	0.057 (0.044)	0.069 (0.067)
C(Incentives)[16]*Diff.	0.067* (0.034)	0.092** (0.042)	0.123** (0.054)	0.040 (0.079)
C(Incentives)[32]*Diff.	0.102** (0.042)	0.121** (0.048)	0.142** (0.067)	0.137 (0.095)
Pseudo R-squared	0.02	0.02	0.04	0.06
No. observations	28760	23480	16080	10360

Clustered standard errors in parentheses.

* $p < .1$, ** $p < .05$, *** $p < .01$

F Allowing for State Dependence

Table 15 presents regression estimates allowing for state dependence—the probability of being correct when there are more 7-sided polygons and separately recording when there are more 9-sided polygons.

Table 15: Allowing for state dependence

	1 minute	2 minutes	5 minutes	12 minutes
Intercept	0.274*** (0.072)	0.295*** (0.079)	0.420*** (0.096)	0.667*** (0.131)
C(state)[9-sided]	-0.030 (0.104)	0.008 (0.108)	-0.123 (0.124)	-0.363** (0.161)
C(Diff.)[2]	0.015 (0.098)	0.091 (0.109)	0.093 (0.141)	0.016 (0.179)
C(Diff.)[3]	0.052 (0.099)	0.186 (0.113)	0.192 (0.149)	0.187 (0.197)
C(Diff.)[6]	0.388*** (0.123)	0.451*** (0.136)	0.435*** (0.164)	0.398* (0.208)
C(state)[9-sided]*C(Diff.)[2]	0.162 (0.144)	0.056 (0.158)	0.183 (0.200)	0.295 (0.251)
C(state)[9-sided]*C(Diff.)[3]	0.226* (0.133)	0.122 (0.142)	0.391** (0.172)	0.572*** (0.220)
C(state)[9-sided]*C(Diff.)[6]	0.327** (0.167)	0.354* (0.182)	0.641*** (0.223)	0.805*** (0.263)
Incentives	0.118*** (0.038)	0.153*** (0.044)	0.220*** (0.059)	0.290*** (0.097)
C(state)[9-sided]*Incentives	0.032 (0.052)	0.010 (0.058)	0.030 (0.075)	0.158 (0.102)
C(Diff.)[2]*Incentives	0.051 (0.055)	0.037 (0.063)	0.065 (0.086)	0.010 (0.123)
C(Diff.)[3]*Incentives	0.017 (0.052)	-0.032 (0.061)	0.060 (0.081)	0.046 (0.128)
C(Diff.)[6]*Incentives	0.157** (0.065)	0.204*** (0.077)	0.278*** (0.106)	0.340** (0.163)
C(state)[9-sided]*C(Diff.)[2]*Incentives	-0.039 (0.077)	0.016 (0.089)	0.020 (0.120)	0.047 (0.155)
C(state)[9-sided]*C(Diff.)[3]*Incentives	0.004 (0.071)	0.100 (0.081)	0.004 (0.111)	-0.049 (0.156)
C(state)[9-sided]*C(Diff.)[6]*Incentives	-0.106 (0.086)	-0.086 (0.100)	-0.186 (0.143)	-0.264 (0.209)
Pseudo R-squared	0.02	0.02	0.04	0.05
No. observations	28760	23480	16080	11400

Clustered standard errors in parentheses.

* $p < .1$, ** $p < .05$, *** $p < .01$

For the easier tasks in which the difference between the numbers of non-decoy shapes is set either to 3 or 6, the state with more 9-sided polygons is easier to identify than the state with more 7-sided polygons.

References

- Jason Abaluck, Jonathan Gruber, and Ashley Swanson. Prescription drug use under Medicare Part D: A linear model of nonlinear budget sets. *Journal of Public Economics*, 164:106–138, 2018.
- George A. Akerlof and Robert J. Shiller. *Phishing for phools: The economics of manipulation and deception*. Princeton University Press, 2015.
- Pedro Bessone, Gautam Rao, Frank Schilbach, Heather Schofield, and Mattie Toma. Sleepless in Chennai: The Consequences of Increasing Sleep among the Urban Poor. Technical report, mimeo, 2019.
- Saurabh Bhargava and Dayanand Manoli. Psychological frictions and the incomplete take-up of social benefits: Evidence from an IRS field experiment. *American Economic Review*, 105(11):3489–3529, 2015.
- Saurabh Bhargava, George Loewenstein, and Justin Sydnor. Choose to lose: Health plan choices from a menu with dominated option. *Quarterly Journal of Economics*, 132(3):1319–1372, 2017.
- David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953.
- Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Saliency theory of choice under risk. *Quarterly Journal of Economics*, 127(3):1243–1285, 2012.
- Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Memory, attention, and choice. *Quarterly Journal of Economics*, Forthcoming, 2020.
- Matthew Botvinick and Wouter Kool. Mental labour. *Nature Human Behaviour*, 2018.
- Charles M. Cameron and Lewis A. Kornhauser. Appeals mechanisms, litigant selection, and the structure of judicial hierarchies. *Institutional Games and the US Supreme Court*, 173:204, 2006.
- Xinyu Cao and Juanjuan Zhang. Preference learning and demand forecast. *Working Paper*, 2019.
- Andrew Caplin and Mark Dean. Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203, 2015.

- Andrew Caplin and Daniel Martin. A testable theory of imperfect perception. *Economic Journal*, 125 (582):184–202, February 2015.
- Andrew Caplin and Daniel Martin. Framing as information design. *Working Paper*, 2019.
- Andrew Caplin, Mark Dean, and John Leahy. Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy. Technical report, National Bureau of Economic Research, 2017.
- Tom Chang, Joshua Graff Zivin, Tal Gross, and Matthew Neidell. Particulate pollution and the productivity of pear packers. *American Economic Journal: Economic Policy*, 8(3):141–69, 2016.
- Tom Y. Chang, Joshua Graff Zivin, Tal Gross, and Matthew Neidell. The effect of pollution on worker productivity: evidence from call center workers in china. *American Economic Journal: Applied Economics*, 11(1):151–72, 2019.
- Raj Chetty, Adam Looney, and Kory Kroft. Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–77, 2009.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, 2006.
- Janet M. Currie. *The invisible safety net: Protecting the nation's poor children and families*. Princeton University Press, 2008.
- Henrique de Oliveira, Tommaso Denti, Maximilian Mihm, and Kemal Ozbek. Rationally inattentive preferences and hidden information costs. *Theoretical Economics*, 12(2):621–654, 2017.
- Joshua T. Dean. Noise, cognitive function, and worker productivity. Technical report, mimeo, 2019.
- Mark Dean and Nathaniel Neligh. Experimental tests of rational inattention. Technical report, Columbia University, 2019.
- Stanislas Dehaene, Hakwan Lau, and Sid Kouider. What is consciousness, and could machines have it? *Science*, 358(6362):486–492, 2017.
- Ambuj Dewan and Nathaniel Neligh. Estimating information cost functions in models of rational inattention. Technical report, Columbia University, 2017.

- Min Ding. An incentive-aligned mechanism for conjoint analysis. *Journal of Marketing Research*, 44(2):214–223, 2007.
- Min Ding, Rajdeep Grewal, and John Liechty. Incentive-aligned conjoint analysis. *Journal of Marketing Research*, 42(1):67–82, 2005.
- Molla S. Donaldson, Janet M. Corrigan, Linda T. Kohn, et al. *To err is human: building a safer health system*, volume 6. National Academies Press, 2000.
- Benjamin Enke. What you see is all there is. *Quarterly Journal of Economics*, Forthcoming, 2020.
- Benjamin Enke and Florian Zimmermann. Correlation neglect in belief formation. *Review of Economic Studies*, 86(1):313–332, 2017.
- Gustav Theodor Fechner. *Elemente der Psychophysik (Elements of Psychophysics)*. Leipzig: Breitkopf und Hartel., 1860.
- Ernst Fehr and Antonio Rangel. Neuroeconomic foundations of economic choice—recent advances. *Journal of Economic Perspectives*, 25(4):3–30, 2011.
- Raymond P. Guiteras and B. Kelsey Jack. Productivity in piece-rate labor markets: Evidence from rural malawi. *Journal of Development Economics*, 131:42–61, 2018.
- Benjamin R. Handel and Jonathan T. Kolstad. Health insurance for” humans”: Information frictions, plan choice, and consumer welfare. *American Economic Review*, 105(8):2449–2500, 2015.
- Benjamin R. Handel and Joshua Schwartzstein. Frictions or mental gaps: what’s behind the information we (don’t) use and when do we care? *Journal of Economic Perspectives*, 32(1):155–78, 2018.
- Rema Hanna, Sendhil Mullainathan, and Joshua Schwartzstein. Learning through noticing: Theory and evidence from a field experiment. *Quarterly Journal of Economics*, 129(3):1311–1353, 2014.
- Justine S. Hastings and Jeffrey M. Weinstein. Information, school choice, and academic achievement: Evidence from two experiments. *Quarterly Journal of Economics*, 123(4):1373–1414, 2008.

- Jiaxiu He, Haoming Liu, and Alberto Salvo. Severe air pollution and labor productivity: Evidence from industrial towns in china. *American Economic Journal: Applied Economics*, 11(1):173–201, 2019.
- William Hoiles, Vikram Krishnamurthy, and Kunal Pattanayak. Rationally inattentive inverse reinforcement learning explains youtube commenting behavior. *arXiv:1910.11703*, 2019.
- Supreet Kaur, Sendhil Mullainathan, Suanna Oh, and Frank Schilbach. Does financial strain lower productivity? *Working Paper*, 2019.
- Jeffrey R. Kling, Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel. Comparison friction: Experimental evidence from medicare drug plans. *Quarterly Journal of Economics*, 127(1):199–235, 2012.
- Jeffrey A. Linder, Jason N. Doctor, Mark W. Friedberg, Harry Reyes Nieva, Caroline Birks, Daniella Meeker, and Craig R. Fox. Time of day and the decision to prescribe antibiotics. *JAMA Internal Medicine*, 174(12):2029–2031, 2014.
- Ulrike Malmendier and Stefan Nagel. Depression babies: do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics*, 126(1):373–416, 2011.
- Filip Matějka and Alisdair McKay. Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. *American Economic Review*, 105(1):272–98, 2015.
- William Morrison and Dmitry Taubinsky. Rules of thumb and attention elasticities: Evidence from under- and overreaction to taxes. 2019.
- Sendhil Mullainathan and Eldar Shafir. *Scarcity: Why having too little means so much*. Times Books – Henry Holt and Co., 2013.
- Sebastian Musslick, Jonathan D. Cohen, and Amitai Shenhav. Estimating the costs of cognitive control from task performance: theoretical validation and potential pitfalls. *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*, 2018.
- Roger Ratcliff. A theory of memory retrieval. *Psychological Review*, 85(2):59, 1978.
- Ralph Tyrell Rockafellar. *Convex Analysis*. Princeton University Press, 1971.

- Heather Schofield. The economic costs of low caloric intake: Evidence from india. Technical report, 2014.
- Joshua Schwartzstein. Selective attention and learning. *Journal of the European Economic Association*, 12(6):1423–1452, 2014.
- Steven Shavell. *Economic analysis of accident law*. Harvard University Press, 2009.
- Christopher A. Sims. Stickiness. *Carnegie-Rochester Conference Series on Public Policy*, 49:317 – 356, 1998.
- Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3): 665–690, 2003.
- Matthew C. Stephenson. Information acquisition and institutional design. *Harvard Law Review*, 124: 1422, 2010.
- George J. Stigler. The economics of information. *Journal of Political Economy*, 69(3):213–225, 1961.
- Dmitry Taubinsky and Alex Rees-Jones. Attention variation and welfare: theory and evidence from a tax salience experiment. *Review of Economic Studies*, 85(4):2462–2496, 2018.
- Shelley E. Taylor and Suzanne C. Thompson. Stalking the elusive “vividness” effect. *Psychological Review*, 89(2):155, 1982.
- Amos Tversky and Daniel Kahneman. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458, 1981.
- Hal R. Varian. *Intermediate Microeconomics: A Modern Approach: Ninth International Student Edition*. WW Norton & Company, 2014.
- Jessica A. Wachter and Michael J. Kahana. A retrieved-context theory of financial decisions. *Working Paper*, 2019.
- Ernst Heinrich Weber. *De Pulsu, Resorptione, Auditu et Tactu. Annotationes Anatomicae et Physiologicae*. C.F. Koehler, Leipzig, Germany, 1834.
- Thomas A. Weber. Simple methods for evaluating and comparing binary experiments. *Theory and Decision*, 69(2):257–288, 2010.