

NBER WORKING PAPER SERIES

RATIONAL INATTENTION, COMPETITIVE SUPPLY, AND PSYCHOMETRICS

Andrew Caplin
Dániel Csaba
John Leahy
Oded Nov

Working Paper 25224
<http://www.nber.org/papers/w25224>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2018

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 Benabou, Mark Dean, Sam Gershman, Wei Ji Ma, Daniel Martin and David Redish for helpful discussions. We thank the Alfred P. Sloan and NOMIS Foundations for support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Andrew Caplin, Dániel Csaba, John Leahy, and Oded Nov. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Rational Inattention, Competitive Supply, and Psychometrics
Andrew Caplin, Dániel Csaba, John Leahy, and Oded Nov
NBER Working Paper No. 25224
November 2018
JEL No. C14,C91,D01,D04,D11,D6,D8

ABSTRACT

Costs of attention, while central to choice behavior, have proven hard to measure. We introduce a simple method of recovering them from choice data. Our recovery method rests on the observation 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. This analogy extends to welfare analysis: consumer welfare net of attention costs is measured in precisely the same way as the profits of a competitive firm. We implement our recovery method in a purpose-built experiment. We quantitatively assess the trade-off between reward level and task complexity. Estimated attention costs are highly correlated with decision time, an important common input in process-based models of attention.

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

Dániel Csaba
New York University
19 West 4th Street
New York, NY 10003
daniel.csaba@nyu.edu

John Leahy
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

Oded Nov
NYU Tandon School of Engineering
6 MetroTech Center
Brooklyn, NY 11201
onov@nyu.edu

1 Introduction

We introduce and implement a simple method of recovering attention costs from choice data. Understanding how these costs impact choice has long been central to economic analysis. Attention costs are as important for demand theory as are costs of production for supply theory. They determine consumer demand and consumer welfare just as a competitive firm’s costs determines its supply curve and its profitability. Yet costs of attention have proven as hard to measure as they are ubiquitous. Our recovery method therefore bridges a key gap between theory and practice.

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 like measuring the output of a firm rather than its profit. Our incentive-based psychometric curve (IPC) 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.

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. We can identify consumer welfare net of costs precisely as one can identify producer surplus from the supply curve of a competitive firm. We can test for the existence of any rationalizing cost function as one can use revealed profitability analysis (e.g. [Varian \[2014\]](#)) to test the theory of profit maximization. We can extract rich qualitative information from each experimental observation by noting that any strategies that achieve higher gross utility must be more subjectively costly. As with profitability, we can also provide a lower bound on the incremental cost of unchosen strategies.

We illustrate our approach in a purpose-built experiment in which we vary both the linear incentive and task difficulty. The incentive is implemented by assigning different probability points to different rounds of the experiment. We implement probabilistic payment using the last digits of the computer’s millisecond timer.

We confirm that the resulting experimental data satisfy simple behavioral tests for the existence of a rationalizing cost function. We estimate these cost functions for each fixed level of difficulty. We find that they can all be fit well by logit models with a common log-transform of the incentive level. With the cost functions in hand we carry out welfare analysis, and assess the trade-off between task complexity and rewards. We quantify the extent to which rewards would be sacrificed in moving from a more challenging to an easier task.

One of our experimental treatments sets all incentives to zero. For that reason we are able to recover not only the IPC, 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 \[1860\]](#)). In that sense our experiment is two dimensional. By varying both the incentive and stimulus strength we produce a psychometric surface that quantifies the trade-off between task complexity and incentives. We introduce iso-performance curves that capture this trade-off visually.

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, energy use, etc. 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 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.

Our approach links with several contemporary lines of psychological research in which individuals are seen as weighing the benefits of cognitive effort against an inherent cost. This includes the resource-rational approach to cognition [[Griffiths et al., 2015](#)], the computationally rational approach [[Gershman et al., 2015](#)], and algorithmic rationality [[Halpern and Pass, 2011](#)]. In fact psychologists have used microeconomic tools at least since [Navon and Gopher \[1979\]](#) (see [Botvinick and Kool \[2018\]](#)). Even the conceptualization of this effort as producing information is essentially identical: “To convert mental effort into an approachable object of scientific study, a useful first step is to operationalize it [...] in terms of information processing...” [[Shenhav et al., 2017](#), p. 3].

[Musslick et al. \[2018\]](#) provided a telling indication of just how strong is the recent convergence of interest between economics and psychology. For essentially identical reasons, they establish a recovery result for a simple class of one dimensional tasks. Their result also rests on the assumption of risk neutrality. Our recovery result applies more generally to all decision problems and is independent of risk aversion.

Attentional effort is far from the only form of subjectively costly mental labor studied by psychologists (see [Griffiths et al. \[2015\]](#) and [Botvinick and Kool \[2018\]](#)). Research on memory and attention is particularly focused on the apparent costs of holding accurate memories of several items in working memory. Traditionally, this has been explained as resulting from a hard constraint on available coding resources (e.g. neural spikes). [van den Berg and Ma \[2017\]](#) develop and test a theory that is based on flexibly trading off behavioral performance against neural costs in the style of RI theory. Our methods may therefore be of broader value. We also advance research in the growing field of human-computer interaction. Ease of use and usefulness of computer interfaces have been studied extensively in this literature, largely based on participants’ own subjective ratings [[Sonderegger et al., 2014](#), [Venkatesh et al., 2003](#), [Zimmerman et al., 2011](#)]. With our methods, these are separately recovered from choice data.

In section 2 we introduce the IPC and the connection with competitive supply. In section 3 we establish

the recovery theorem and relate it to earlier recovery results (e.g. [de Oliveira et al. \[2017\]](#) and [Caplin et al. \[2017\]](#)) and to the recent work of [Musslick et al. \[2018\]](#). In section 4 we illustrate the value of linking attention theory to elementary microeconomics. The experimental design is in section 5. The experimental results are in section 6. Section 7 concludes. The Appendices include experimental instructions and experimental robustness checks.

2 The IPC

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.

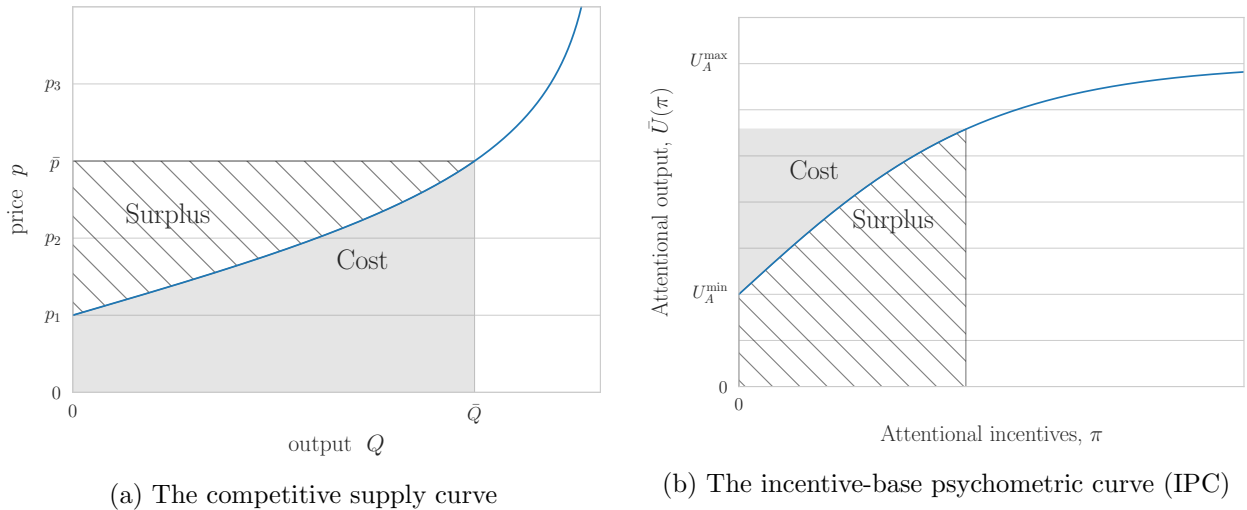


Figure 1: Analogy with competitive supply

2.1 Attentional Incentives, π

We introduce the general attentional problem and corresponding linear incentives that mimic the role of output price in competitive supply. The attentional incentive is presented on the horizontal axis in the IPC graph of figure 1b.

A decision maker (DM) is facing a problem with finitely many possible states, $\omega \in \Omega$. Prior beliefs about the states are denoted by $\mu \in \Delta(\Omega)$. A decision problem comprises a finite set A of $N \geq 2$ actions which yield utility $u(a, \omega) \in \mathbb{R}$ to the DM when choosing action a in state ω . No action is optimal in all states, so that the DM has an incentive to learn the state.

We take advantage of the linearity of expected utility theory in constructing a 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,*

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

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.

2.2 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 [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,

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

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.2 we show that this is a property of rational inattention (RI) models.

2.3 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 $\bar{U}(1)$,

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

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

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 a 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 fixed underlying 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, 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 curve 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 fixed even as the decision problem varies. In essence, RI theory involves fixing a technology of attention. 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.

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 arbitrary decision problem A_π solves the following optimization problem,

$$\max_{P_\pi \in \mathcal{P}(A_\pi)} \sum_{\omega} \sum_n u(a_\pi(n), \omega) P_\pi(n, \omega) - K(P_\pi). \quad (5)$$

Assuming the solution exists we denote any maximizing SDSC strategy as \hat{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

¹Since the prior is held fixed, one can map the joint distributions to a collection of conditional distributions corresponding to general statistical experiments. This allows us to apply notions defined for experiments such as Blackwell informativeness (for a precise definition see appendix A.1) accordingly. Following Blackwell [1953] it is standard that one statistical experiment is more informative than another if it yields a higher ex ante utility in any given decision problem.

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

$$U(P) := \sum_{\omega} \sum_{a \in A(P)} u(a, \omega) P(a, \omega). \quad (6)$$

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

$$\bar{K}_A(u) := \inf_{\{P \in \mathcal{P}(A) | U(P) \geq u\}} K(P). \quad (7)$$

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.2.

The above construction is familiar in information theory. The UCC 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,

$$Q(p) := \arg \max_{Q \geq 0} pQ - C(Q). \quad (8)$$

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 UCC we can write the DM's problem (5) for any $\pi > 0$ as,

$$\max_{u \in [U_A^{\min}, U_A^{\max}]} \pi u - \bar{K}_A(u). \quad (9)$$

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 9, $\hat{u}(\pi)$, 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, $\pi \bar{U}(\pi) = \pi \hat{u}(\pi)$. Hence, the IPC plots the maximizer of problem 9 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 an RI DM can be recovered from the IPC as,*

$$\bar{K}_A(\bar{U}(\pi)) = \pi \bar{U}(\pi) - \int_0^\pi \bar{U}(t) dt. \quad (10)$$

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 UCC is well-behaved. We use results in convex analysis [Rockafellar, 1971] relating the IPC to the convex conjugate of the UCC to prove the theorem in appendix A.3.²

²Note that the behavioral data generated by a UCC and its convex hull are indistinguishable from one another and




Figure 2: Cost recovery and the IPC

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](#)].

We demonstrate and implement our approach in sections 5 - 6.

hence from data one can only recover the UCC 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 UCC is weakly increasing in the Blackwell order.

4 Applications

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” (NIAS) [Caplin and Martin, 2015], which insists on Bayesian expected utility maximization. The other condition, “No Improving Attention Cycles” (NIAC) [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 (NIAS): *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 (NIAC): *Given a finite sequence of decision problems and corresponding SDSC data sets,*

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

with $A_1 = A_M$,

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

Dean and Neligh [2017] introduce statistical tests of these conditions which we apply in our experimental data.

4.2 Welfare

When the NIAS and NIAC tests are passed, 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 .

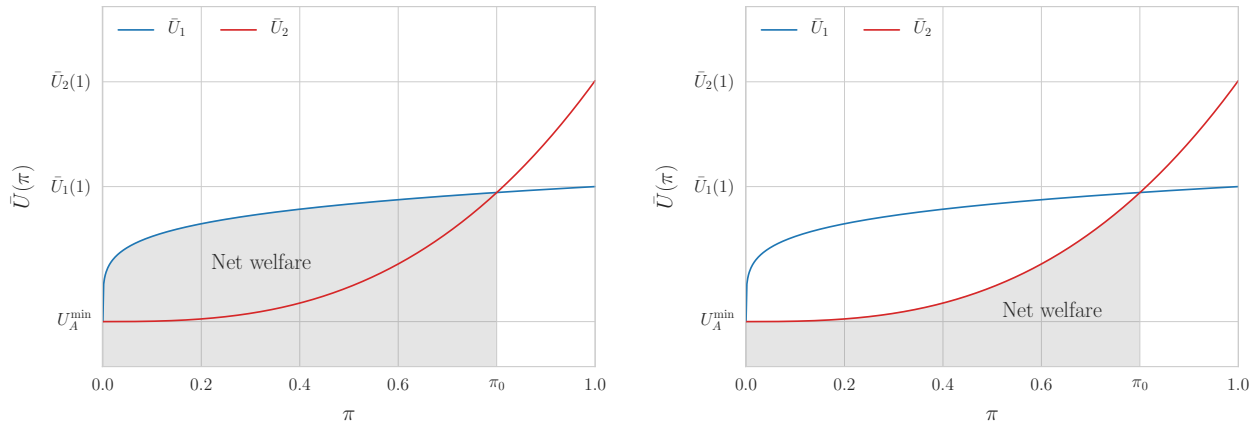


Figure 3: Net welfare

4.3 More Informative and Revealed More Costly

Based on the analogy with elementary production theory and revealed profitability [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 5 (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.³ 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,

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

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,

³In the binary case one experiment is more Blackwell informative than another if it lies to the north-east of it [Weber, 2010, p. 269].

	$a_\pi(1)$	$a_\pi(2)$
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 UCC,

$$\bar{K}_A(u) = u \log(u) + (1 - u) \log(1 - u) - \log(0.5). \quad (11)$$

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 indeed precisely the UCC associated with the Shannon cost in [11](#).

5 Experimental Design

Our experimental design implements the theoretical framework while addressing a number of key practical considerations.

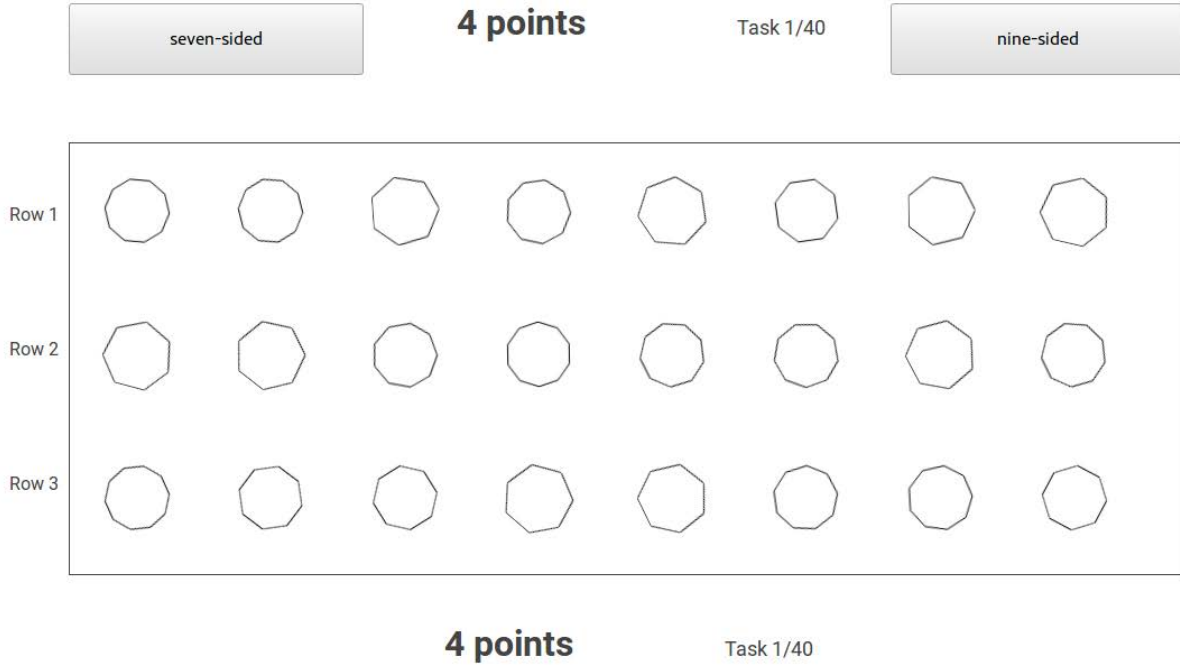


Figure 5: Example of an experimental task

As illustrated in figure 5 the 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.

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 any of the non-decoy shapes being more common has equal chance 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 numbers of rounds given incentive levels are 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 averaging in this form 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 Attentive Subjects

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. Our best proxy for inattentive behavior is the total time spent on the 40 rounds. In recovering the cost function

from experimental data we consider only subjects who spend at least 5 minutes in total on the whole experiment. This threshold leaves us with 402 subjects out of the initial 814. In appendix C.1 we detail how we set this threshold and conduct sensitivity analyses.

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 both the attentive and inattentive subsamples.




Figure 7: Average precision across rounds

As seen in figure 7 we find that there is no significant variation in performance across time for either the attentive or inattentive subjects. As expected the average performance levels are significantly different for attentive and inattentive subsamples. We detail the analysis in appendix C.2.

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,

$$\begin{aligned}
 u_\pi(a(7), \omega_7) &= u_\pi(a(9), \omega_9) = \pi u(\$10), \\
 u_\pi(a(9), \omega_7) &= u_\pi(a(7), \omega_9) = 0.
 \end{aligned}$$

The IPC is therefore,

$$\begin{aligned}\bar{U}(\pi) &= \left(\mu(\omega_7)P_\pi(a(7) \mid \omega_7) + \mu(\omega_9)P_\pi(a(9) \mid \omega_9) \right) u(\$10) \\ &=: \bar{P}_\pi u(\$10).\end{aligned}\tag{12}$$

Equation 12 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 NIAS conditions collapse to,

$$P_\pi(a(7) \mid \omega_7) - P_\pi(a(7) \mid \omega_9) \geq 0;\tag{13}$$

$$P_\pi(a(9) \mid \omega_9) - P_\pi(a(9) \mid \omega_7) \geq 0.\tag{14}$$

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 NIAS condition to check. Four difficulty levels and seven incentive levels leave us with $4 \cdot 7 = 28$ instances to check for the NIAS inequality. When the incentive level is $\pi = 0$ the NIAS condition is vacuous, since without incentives rational subjects could pick arbitrarily. This means that effectively there are 24 non-vacuous NIAS conditions to check. We follow [Dean and Neligh \[2017\]](#) and compute the empirical counterparts of the inequalities to formulate statistical tests. We report p-values for both the attentive and inattentive samples.

π	Diff.			
	1	2	3	6
0	.018	.012	.448	.001
1	.311	.598	.010	.000
2	.089	.024	.000	.000
4	.028	.003	.015	.000
8	.003	.245	.030	.007
16	.069	.007	.005	.000
32	.837	.095	.033	.000

Table 2: NIAS test p-values
Inattentive sample

π	Diff.			
	1	2	3	6
0	.008	.000	.000	.000
1	.000	.000	.000	.000
2	.000	.000	.000	.000
4	.000	.000	.000	.000
8	.000	.000	.000	.000
16	.000	.000	.000	.000
32	.000	.000	.000	.000

Table 3: NIAS test p-values
Attentive sample

Considering the attentive sample, all 24 non-vacuous NIAS tests are passed at the 1% significance level. In the inattentive sample only 12 out of the 24 NIAS tests are passed at the 1% significance level.

Given the symmetry of the problem if the NIAC conditions hold, the NIAC conditions imply 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 NIAC conditions for one difficulty level at a time.

We report the increment in average probabilities for the inattentive and attentive subsamples below. A positive increment implies that the corresponding NIAC condition is satisfied.

$\Delta\pi$ \ Diff.	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 4: Inattentive sample

$\Delta\pi$ \ Diff.	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 5: Attentive sample

We check the adjacent pairwise comparisons for each difficulty level separately. This leaves us with 24 tests to check for the NIAC condition. As expected we can see that there are fewer violations of the NIAC condition in the attentive sample than the inattentive sample. In the attentive sample 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, in the inattentive sample there are 10 violations altogether. We report the significance level of these tests in the on-line appendix.

Altogether, we confirm that the behavior revealed by the attentive subsample 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 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




Figure 8: IPC for the simplest task

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.⁴ This is illustrated in figure 8.

We estimate the IPCs across all difficulty levels jointly. We find that a logit model with a log-transform of the geometric incentive scheme provides a relatively good fit that also lends itself to interpretability and easy calculations for costs and net welfare. 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 ,

$$\log\left(\frac{\bar{P}_i}{1 - \bar{P}_i}\right) = \alpha + \sum_{\text{Diff} \in \{2,3,6\}} \alpha_{\text{Diff}} D_{\text{Diff},i} + \beta * \log(\pi_i + 1) + \sum_{\text{Diff} \in \{2,3,6\}} \beta_{\text{Diff}} D_{\text{Diff},i} \log(\pi_i + 1), \quad (15)$$

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 $\text{Diff} = 1$. 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.

⁴As 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 E 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.




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 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 attentive sample included in the estimation. We provide regression tables showing the sensitivity analysis and joint significant tests in appendix D.

An interesting feature is that the differences across difficulty levels are relatively stable across incentive levels, implying that one can learn most of the differences by just looking at the intercept—the slopes being similar.

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.

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.




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

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 6 shows the corresponding incentive levels, achieved probabilities of being correct, costs, and net welfare. Cost and net welfare are expressed as percentages of $u(\$10)$.

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

The patterns are clear—the cost of achieving a level of gross utility is monotonically increasing in 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 (SPC) varying the difference between the number of non-decoy shapes. Figure 11 plots this traditional non-incentivized SPC using our experimental data.




Figure 11: The SPC with no incentive

Figure 12 plots the SPCs as a function of discriminability for each incentive level. The lowest SPC 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 SPC higher.




Figure 12: SPCs 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 between these two important features of a task through simple iso-performance curves as depicted in figure 13b.

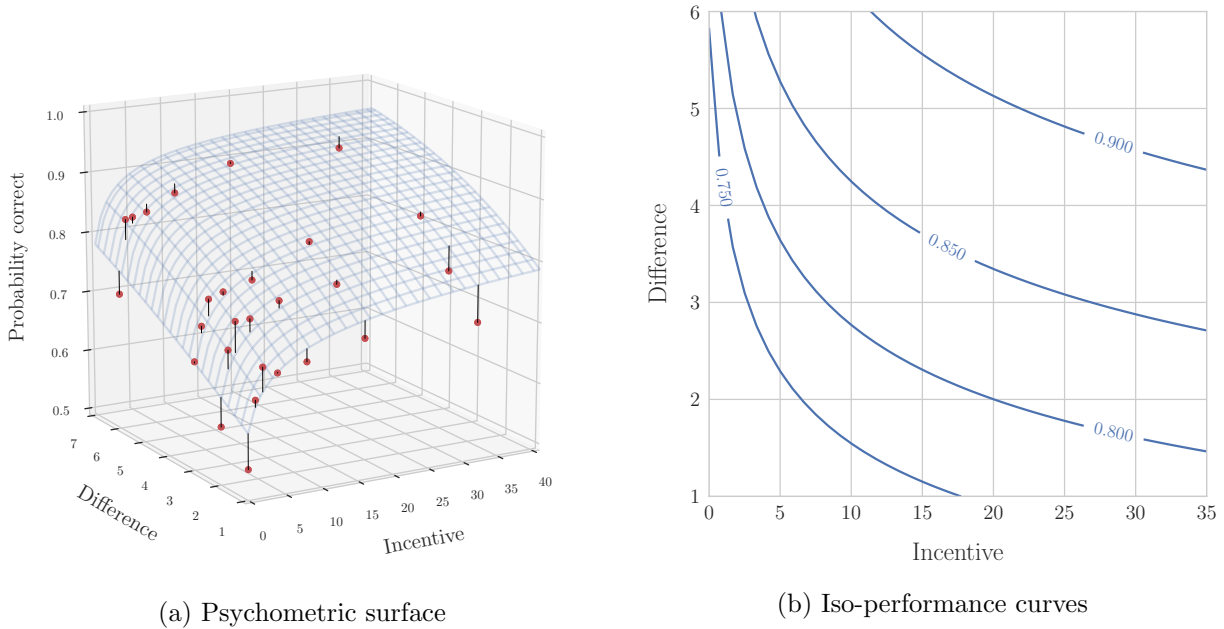


Figure 13: Trade-off between incentive and task complexity

As we see in the graph, if task complexity increases the DM has to be compensated with higher incentives to stay on 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.⁵

⁵Since the probabilities of being correct in the data do not lie exactly on the IPC we use the predicted value from the IPC directly.

Table 7 shows the correlations across all these variables.

	probability correct	mean time	median time	cost from IPC
probability correct	1.000	0.436	0.560	0.412
mean time	0.436	1.000	0.933	0.797
median time	0.560	0.933	1.000	0.880
cost from IPC	0.412	0.797	0.880	1.000

Table 7: 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.

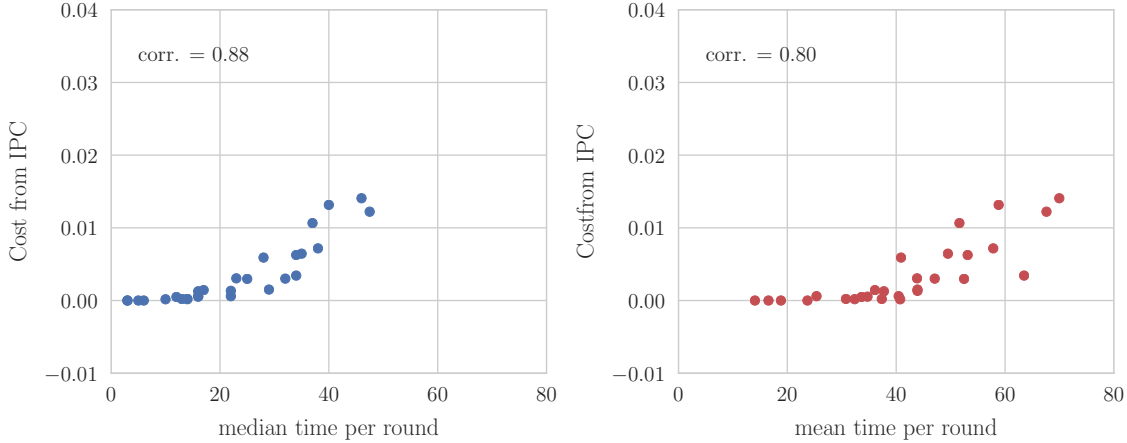


Figure 14: Correlation between estimated cost and time taken

7 Concluding Remarks

We introduce and implement a simple method of recovering attention costs from choice data. Our IPC 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 experimentally implement our recovery method in tasks of varying complexity. The data pass tests for applicability of our approach. We quantitatively assess the welfare trade-off between rewards and task complexity. We relate the costs recovered from choice data to decision time, an important psychological input.

Andrew Caplin

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

Dániel Csaba

Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
daniel.csaba@nyu.edu

John Leahy

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

Oded Nov

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

References

- David Blackwell. Equivalent comparisons of experiments. *Ann. Math. Statist.*, 24(2):265–272, 06 1953. doi: 10.1214/aoms/1177729032. URL <http://dx.doi.org/10.1214/aoms/1177729032>.
- Matthew Botvinick and Wouter Kool. Mental labour. *Nature Human Behaviour*, 2018.
- 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. *The Economic Journal*, 125(582):184–202, February 2015.
- Andrew Caplin, Mark Dean, and John Leahy. Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy. Technical report, National Bureau of Economic Research, 2017.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, 2006.
- Henrique de Oliveira, Tommaso Denti, Maximilian Mihm, and Kemal Ozbek. Rationally inattentive preferences and hidden information costs. *Theoretical Economics*, 12(2):621–654, 2017.
- Mark Dean and Nathaniel Neligh. Experimental tests of rational inattention. Technical report, Columbia University, 2017.
- Ambuj Dewan and Nathaniel Neligh. Estimating information cost functions in models of rational inattention. Technical report, Columbia University, 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.
- Samuel J. Gershman, Eric J. Horvitz, and Joshua B. Tenenbaum. Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245):273–278, 2015.
- Thomas L. Griffiths, Falk Lieder, and Noah D. Goodman. Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7(2):217–229, 2015.
- Joseph Y. Halpern and Rafael Pass. Algorithmic rationality: Adding cost of computation to game theory. *ACM SIGecom Exchanges*, 10(2):9–15, 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.

- 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.
- David Navon and Daniel Gopher. On the economy of the human-processing system. *Psychological review*, 86(3):214, 1979.
- Roger Ratcliff. A theory of memory retrieval. *Psychological Review*, 85(2):59, 1978.
- Ralph Tyrell Rockafellar. *Convex Analysis*. Princeton University Press, 1971.
- Amitai Shenhav, Sebastian Musslick, Falk Lieder, Wouter Kool, Thomas L. Griffiths, Jonathan D. Cohen, and Matthew M. Botvinick. Toward a rational and mechanistic account of mental effort. *Annual review of neuroscience*, 40:99–124, 2017.
- 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.
- Andreas Sonderegger, Andreas Uebelbacher, Manuela Pugliese, and Juergen Sauer. The influence of aesthetics in usability testing: the case of dual-domain products. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 21–30. ACM, 2014.
- Ronald van den Berg and Wei Ji Ma. A rational theory of the limitations of working memory and attention. *bioRxiv*, page 151365, 2017.
- Hal R. Varian. *Intermediate Microeconomics: A Modern Approach: Ninth International Student Edition*. WW Norton & Company, 2014.
- Viswanath Venkatesh, Michael G. Morris, Gordon B. Davis, and Fred D. Davis. User acceptance of information technology: Toward a unified view. *MIS quarterly*, pages 425–478, 2003.
- 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.
- John Zimmerman, Anthony Tomasic, Charles Garrod, Daisy Yoo, Chaya Hiruncharoenvate, Rafae Aziz, Nikhil Ravi Thiruvengadam, Yun Huang, and Aaron Steinfeld. Field trial of tiramisu: crowdsourcing bus arrival times to spur co-design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1677–1686. ACM, 2011.

A Proofs

A.1 Definition of Blackwell Informativeness

Below we define the notion of Blackwell informativeness, a concept widely used in economics, on the SDSC strategy space that we use.

Definition 6. For any $P \in \mathcal{P}$ define the marginal probability of action, a , and the revealed posterior at any action, a ,

$$P(a) := \sum_{\omega} P(a, \omega) \quad \forall a \in A(P); \quad (16)$$

$$\gamma_P^a(\omega) := \frac{P(a, \omega)}{P(a)} \quad \forall (a, \omega) \in A(P) \times \Omega. \quad (17)$$

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,

$$\sum_{a \in A(P)} \phi(\gamma_P^a) P(a) \geq \sum_{a \in A(Q)} \phi(\gamma_Q^a) Q(a), \quad (18)$$

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

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.

A.2 Blackwell Monotonicity and the UCC

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

Proof.

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 a $t \in [0, 1]$ such that

$$U(tP_2 + (1-t)P^I) = u_1,$$

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

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

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

A.3 Proof of Theorem 1

For any $\pi > 0$ define,

$$\widehat{N}(\pi) := \max_{P \in \mathcal{P}_\pi} \sum_{\omega} \sum_n u(a_\pi(n), \omega) P(n, \omega) - K(P). \quad (19)$$

Consider an arbitrary $\pi_0 > 0$. This is associated with net utility $\widehat{N}(\pi_0)$ and normalized expected utility $\bar{U}(\pi_0)$. Consider now any other $\pi > 0$. The optimal strategy at π must do at least as well as applying \widehat{P}_{π_0} ;

$$\begin{aligned} \widehat{N}(\pi) &\geq \sum_{\omega} \sum_n \widehat{P}_{\pi_0}(a, \omega) u(a_\pi(n), \omega) - K(\widehat{P}_{\pi_0}) \\ &= \pi \frac{U(\widehat{P}_{\pi_0})}{\pi_0} - K(\widehat{P}_{\pi_0}) \\ &= \pi \bar{U}(\pi_0) - \pi_0 \bar{U}(\pi_0) + \pi_0 \bar{U}(\pi_0) - K(\widehat{P}_{\pi_0}) \\ &= (\pi - \pi_0) \bar{U}(\pi_0) + \widehat{N}(\pi_0); \end{aligned} \quad (20)$$

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 $\widehat{N}(\pi_0)$ and $\bar{U}(\pi_0)$. It follows immediately that \widehat{N} is convex and that $\bar{U}(\pi)$ is an element of the subdifferential of \widehat{N} at point π ,

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

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

$$\widehat{N}(\pi) = \int_0^\pi \bar{U}(t) dt + C, \quad (21)$$

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

To pin down C , note that,

$$\begin{aligned} \lim_{\pi \searrow 0} \widehat{N}(\pi) &= \lim_{\pi \searrow 0} \left\{ U(\widehat{P}_\pi) - K(\widehat{P}_\pi) \right\} = - \lim_{\pi \searrow 0} K(\widehat{P}_\pi), \\ \text{and} \quad \lim_{\pi \searrow 0} \widehat{N}(\pi) &= \lim_{\pi \searrow 0} \int_0^\pi \bar{U}(t) dt + C = C. \end{aligned}$$

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

The definition of the IPC together with equation 21 and C being zero gives

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

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

$$\hat{N}(\pi) = U(\pi) - K(\hat{P}_\pi).$$

Rearranging and using equation 3 gives

$$K(\hat{P}_\pi) = U(\hat{P}_\pi) - \hat{N}(\pi) = \pi \bar{U}(\pi) - \hat{N}(\pi),$$

which is the area above the IPC. □

B Experimental Design

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)

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

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

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

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

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

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

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


They sum to 24 as per algorithm.

B.2 Instructions

The experiment is available at <http://distinguishshapes.herokuapp.com/welcome>. Below we repeat the information found on the website.

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 32 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,

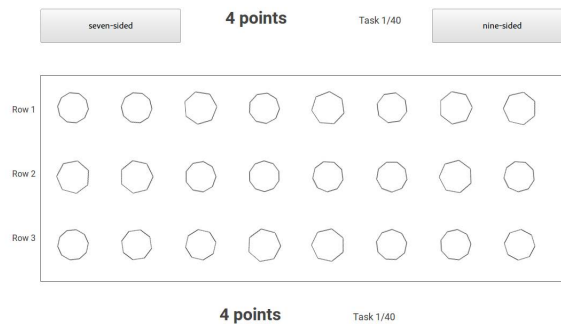
and match your probability of winning to your achieved score: the higher your score is, the more likely you are to win the \$10 bonus. To confirm this randomness, you can start and stop the clock as many times as you want, to check if you can stop it at a number of your choice. Please stop the clock at least 3 times.

18 :43:26:948



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.



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 8 shows the number of subjects for each difficulty level.⁶

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

Table 8: Number of subjects

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

C.1 Time Patterns and Types of Subjects

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

⁶Subjects who started but did not finish the experiment were dropped. There were 24 of them.

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. The distribution of the logarithm of total times is indicative of a bimodal distribution. 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 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.



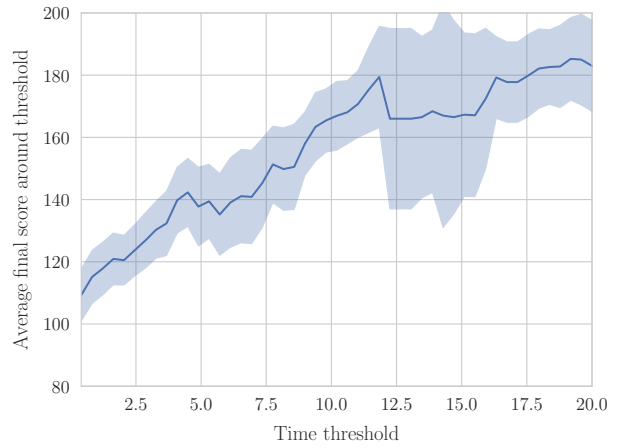


Figure 16: Histogram of final scores (<1 minute of total time)

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 inattentive subjects from attentive ones while retaining power. We 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 D.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 9: Number of subjects spending at least 5 minutes on the experiment

The performance of the 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 the inattentive subsample. On the other hand, attentive subjects show considerable variation as shown in figure 19.




Figure 19: Median times per round for each incentive level

C.2 Learning and Fatigue

There is no significant change in average performance in the order the tasks are presented pooling across all incentive levels. The coefficient on the round rank is not significantly different from zero in both attentive and inattentive subsamples.

Table 10: Performance across rounds

	Attentive	Inattentive
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 in both the attentive and inattentive subsamples show a clear pattern of spending less time per round in later rounds.

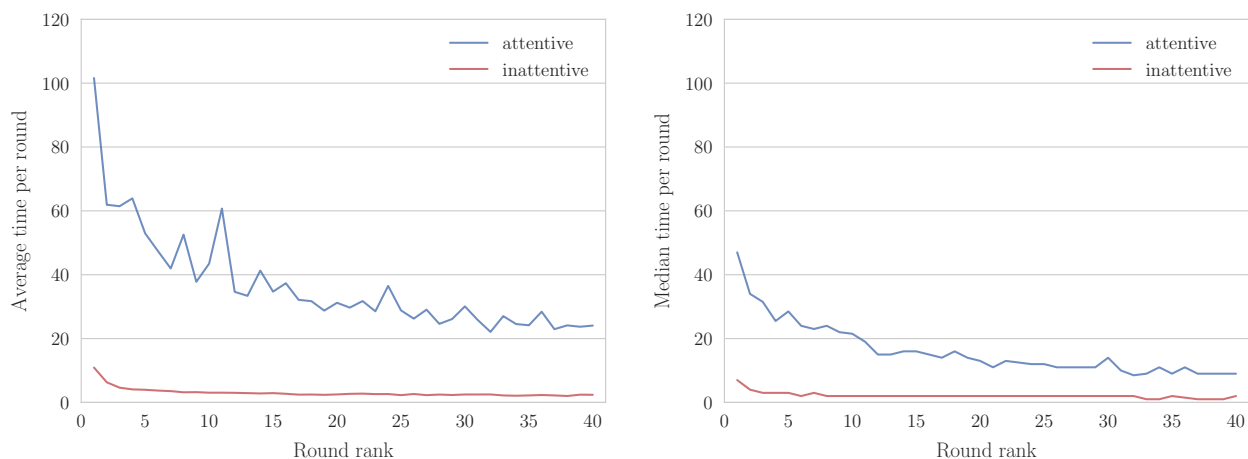


Figure 20: Average and median time across round ranks

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 the 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 confirms that stopping the clock as a randomization device worked well.

D Regression Results

D.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)

D.2 Regression Results for the SPC

Table 13: Regression results for the SPC

	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.

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

E Allowing for State Dependence

Table 14 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 14: 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.