

Redeeming Falsifiability?

Mark Whitmeyer Kun Zhang*

March 20, 2023

Abstract

We revisit Popper's falsifiability criterion. A tester hires a potential expert to produce a theory, offering payments contingent on the observed performance of the theory. We argue that if the informed expert can acquire additional information, falsifiability *does* have the power to identify worthless theories.

*Arizona State University. Emails: mark.whitmeyer@gmail.com & kunzhang@asu.edu.
Comments by Ashwin Kambhampati, Andreas Kleiner, and Joseph Whitmeyer helped us improve this paper immensely.

1 Falsifiability

The criterion of falsifiability—that a scientific idea is one that can be falsified, i.e., conclusively rejected by the data—is central to science. Worryingly, [Olszewski and Sandroni \(2011\)](#) illustrate a flaw with this notion.

They study a model in which a principal (Alice) hires an expert (Bob) to deliver a falsifiable theory, which Alice then checks against a sequence of data points. In order to provide incentives to Bob, Alice uses the carrot (paying Bob a lump sum if he delivers a falsifiable theory) and the stick (fining Bob if his theory is falsified). Even though i) Alice has an unbounded dataset with respect to which she may evaluate the theory, ii) the fine levied on Bob can be unboundedly large, and iii) Bob evaluates his future prospects pessimistically (by the minimum expected utility given any future realization of the data); [Olszewski and Sandroni](#) show that Alice cannot identify an informed Bob from one who is uninformed.

The purpose of this paper is to reexamine this result by allowing the informed Bob to acquire additional information. That is, we endogenize Bob’s knowledge of the data-generating process. We show that as long as the informed Bob can acquire information, the falsifiability criterion *does* “separate the wheat from the chaff:” it distinguishes between the informed and the uninformed Bob.

1.1 Related Literature

There is a large literature studying how clueless agents can evade detection by empirical tests. This collection of papers includes the seminal “calibration” result of [Foster and Vohra \(1998\)](#) (see also [Dawid \(1982\)](#), [Dawid \(1985\)](#), and [Foster \(1999\)](#)), which was followed by [Fudenberg and Levine](#)

(1999), Sandroni et al. (2003) and Hart (2022). Other papers studying the use of tests to catch masquerading non-experts—Lehrer (2001), Dekel and Feinberg (2006), Shmaya (2008), Olszewski and Sandroni (2007), Olszewski and Sandroni (2008), Olszewski and Sandroni (2009), and Hu and Shmaya (2013) (to name a few)—followed over the next few decades.¹

More recently, a sizeable collection of papers exploring rational inattention and endogenous flexible information acquisition has emerged. A subclass of these are those papers that study the contracting problem of paying an agent to acquire information. Rappoport and Somma (2017) study the problem of inducing an agent to acquire hard evidence; and Yoder (2022) studies the impact of private information in this setting, introducing a screening element. Sharma et al. (2020) asks how to impel a risk-neutral agent to acquire and report honestly soft information in a two-state environment;² and Whitmeyer and Zhang (2022) study the general problem of “buying opinions,” allowing for an arbitrary number of states and risk aversion. In both Zermeño (2011) and Clark and Reggiani (2021) both learning and decision making are delegated to an agent.

This paper connects the expert-testing and information-acquisition literatures, allowing the competent expert to acquire information in a classical setting of designing tests to fail charlatans.

¹Olszewski (2015) provides a helpful survey.

²Müller-Itten et al. (2021) formulate an “ignorance equivalent,” which they apply to (among other things) incentivizing an agent to acquire and report information.

2 The Main Result

There is an unknown state of the world $\theta \in \Theta$, where $|\Theta| = n < \infty$. Alice offers a contract to a self-proclaimed expert, Bob, that consists of a lump-sum payment (in utils), $u > 0$, and a penalty, $d > 0$. Henceforth, we refer to this variety of contract as simply a **contract**. If Bob accepts the contract he obtains the payment, u , up front; then announces a state that will not realize. The state subsequently realizes and is publicly observed by Alice and Bob; and if the announced state coincides with the realized state, Alice levies the penalty, d , on Bob.

The expert, Bob, is either informed or uninformed. If Bob is informed, he has a prior $\mu \in \Delta := \Delta(\Theta)$. After seeing the contract, the informed Bob decides whether to acquire additional information at a cost. We can allow him to learn either before or after deciding whether to accept the contract: our main result holds under both specifications. He learns by observing the outcome of a statistical **experiment**, which is a pair $E = (S, \chi)$, where S is a set of possible signals and $\chi: \Theta \rightarrow \Delta(S)$ is a stochastic map from the set of states to the set of signals. The cost of acquiring information is given by a cost functional $C: \mathcal{E} \times \Delta \rightarrow \mathbb{R}_+ \cup \{\infty\}$, where $C(E, \mu)$ is the cost of acquiring experiment E when Bob's prior is μ .

Let

$$\Upsilon(E, \mu) := \min \left\{ \min_i \mu_i, 1 - \mu_i \right\} - \int_S \min \left\{ \min_i \{ \mu_i P_i(s) \}, (1 - \mu_i) P_n(s) \right\} ds,$$

where μ_i is the i -th entry of μ ($i = 1, \dots, n-1$), and $P_i = \chi(\theta_i)$. $d \cdot \Upsilon(E, \mu)$ is the benefit from learning according to experiment E for informed Bob with prior μ . For fixed $\delta \geq 0$, say that an experiment E is **δ -valuable** for prior μ if $\Upsilon(E, \mu) > \delta$. We make the following joint assumption on C and \mathcal{E} :

(A) There exist $\varepsilon, \eta > 0$ such that for all $\mu \in B\left(\frac{1}{\mathbf{n}}, \eta\right)$, there exists $E_\mu \in \mathcal{E}$ that is ε -valuable for μ , and $C(E_\mu, \mu) \leq T$ for some $T \in \mathbb{R}_+$.³

A stronger condition than Assumption (A) is that there is a non-trivial (informative) experiment that informed Bob, regardless of his prior, has access to at a finite cost.⁴ Another example that satisfies Assumption (A) is that informed Bob has access to all experiments, and his cost of acquiring information is posterior separable (Caplin et al. (2022)).⁵ His cost could also be experimental (Denti et al. (2022a)).

If Bob is uninformed he cannot acquire any information. Moreover, the uninformed Bob evaluates payoffs by his minimum expected utility in all possible states. By rejecting the contract, both varieties of Bob get payoff 0.⁶

We say a contract **screens** the uninformed Bob if the informed Bob prefers to accept it but the uninformed Bob strictly prefers to reject. Our main theorem is simple:

Theorem 2.1. *There exists a contract that screens the uninformed Bob.*

³ $\frac{1}{\mathbf{n}}$ denotes the vector that has $\frac{1}{n}$ for each of its entries, and $B\left(\frac{1}{\mathbf{n}}, \eta\right)$ is the ball centered at $\frac{1}{\mathbf{n}}$ with radius η . Note that ε and T are *uniform* to all $\mu \in B\left(\frac{1}{\mathbf{n}}, \eta\right)$.

⁴As a simple example, suppose there are two states, H and L . By incurring a cost $c > 0$, informed Bob can get a signal $s \in \{h, l\}$ such that $\mathbb{P}(h|H) = \mathbb{P}(l|L) = q$, where $1/2 < q \leq 1$.

⁵This class of information costs includes the entropy-based cost function (see e.g. Sims (1998, 2003), and Matějka and McKay (2015)); the log-likelihood cost of Pomatto et al. (2020); the neighborhood-based cost function studied by Hébert and Woodford (2021); and the quadratic (posterior variance) cost function. Some variants of this class of costs can also be allowed; see, for example, Example 4 and 6 in Lipnowski and Ravid (2023).

⁶The perspicacious reader may note that this does not is not a perfect analog of Olszewski and Sandroni's framework. We modified things slightly for the sake of presentation. §2.3 contains a model that is an exact analog of theirs. There, we argue that our insights persist.

Proof. By Assumption (A), there exist $\varepsilon, \eta > 0$ such that for all $\mu \in B\left(\frac{1}{n}, \eta\right)$ we can find $E_\mu \in \mathcal{E}$ such that E_μ is ε -valuable for μ with $C(E_\mu, \mu) \leq T$. If E_μ is ε -valuable for μ , by choosing $d > T/\varepsilon$, $d \cdot \Upsilon(E_\mu, \mu)$ is strictly larger than T for all $\mu \in B\left(\frac{1}{n}, \eta\right)$. Therefore, u can be chosen so that $u - \frac{d}{n} < 0$ but every informed Bob's payoff, net of possible learning costs, is positive. Hence, he accepts the contract. Since the uninformed Bob evaluates payoffs by his minimum expected utility, there is no way of randomizing over announcements that secures him a payoff greater than $u - \frac{d}{n}$. As this is strictly less than 0, he refuses the contract. ■

The remainder of the paper proceeds as follows. In §2.1, we explain how our model maps to the motivating example in Olszewski and Sandroni (2011) and why Theorem 2.1 fails to hold there. In §2.2 we illustrate, via two examples how can we overcome the difficulty by allowing the informed agent to acquire further information. §2.3 argues that our simplification of Olszewski and Sandroni (2011) is innocuous; and §2.4 studies a variant in which the uninformed Bob is also an expected-utility maximizer, but incapable of learning.

2.1 Olszewski and Sandroni's Example, Reframed

Our model (if we took away the informed Bob's ability to acquire information) maps to Olszewski and Sandroni's example as follows. There, the informed Bob knows the composition of an urn that contains balls of n possible colors. The uninformed Bob does not. Bob announces a falsifiable theory: he must claim that at least one color is impossible. Alice then draws a ball, and if it is the "impossible" color, Bob gets fined d .

In our framework, the realized state, θ , corresponds to the color of the

ball drawn by Alice. The informed Bob, therefore, has a prior $\mu \in \Delta$, according to which θ realizes. Bob's announcement is of a state (or states) he claims will not realize.

Without learning, the following result summarizes [Olszewski and Sandroni \(2011\)](#)'s example:

Proposition 2.2. *If a contract is such that an informed Bob accepts it, no matter his prior, an uninformed Bob will also accept it.*

Proof. Let $\mu := (\mu_1, \dots, \mu_{n-1})$ denote an arbitrary prior in Δ . Observe that an informed Bob accepts a contract if $u - d \min_i \mu_i \geq 0$. As $\min_i \mu_i$ is maximized when $\mu_i = \frac{1}{n}$ for all i , an informed Bob accepts a contract, no matter his prior, if and only if $u \geq \frac{d}{n}$. However, as observed by [Olszewski and Sandroni \(2011\)](#), by randomizing uniformly over announcements, the uninformed Bob guarantees himself a payoff of $u - \frac{d}{n}$, which is weakly greater than 0 by construction. Therefore, the uninformed Bob also accepts the contract. ■

In fact, even more general contracts (beyond the proposed one corresponding to “falsifiability”) cannot weed out the ignorant Bob. That is, there is no contract—i.e., a pair (M, t) , where M is a compact set of messages and $t: M \times \Theta \rightarrow \mathbb{R}$ is continuous—such that an informed Bob accepts it (no matter his prior) but an uninformed Bob does not.

For simplicity, suppose there are two states. Observe that a contract induces a convex value function on the 1-simplex, $V(x)$, for the informed Bob. That he accepts it, no matter his prior, requires that $V(x) \geq 0$ for all x . But then it is easy to see that the uninformed Bob will also accept the contract: if $V(0) = 0$ or $V(1) = 0$, the uninformed Bob will send a message that is optimal for the informed Bob at prior 0 or 1, respectively. If $V(0), V(1) > 0$ and $V(\tilde{x}) = 0$ for some $\tilde{x} \in (0, 1)$ for which $V'(\tilde{x})$ exists,

the uninformed Bob will send the message optimal for the informed Bob at prior \tilde{x} , which necessarily has a state-independent payoff of 0. Finally, if $V(0), V(1) > 0$ and $V(\tilde{x}) = 0$ for some $\tilde{x} \in (0, 1)$ but $V'(x) \neq 0$ for all $x \in [0, 1]$ (and so, necessarily V is kinked at \tilde{x}), the uninformed Bob will mix between the two messages optimal for the informed Bob with prior \tilde{x} , and nature will choose the worst-case probability \tilde{x} for the uninformed Bob, yielding him a payoff of 0.

Some Intuition

The basic intuition behind this subsection's finding—and unless we are confused, [Olszewski and Sandroni \(2011\)](#)—is that in order to elicit information from an informed expert, no matter his prior, the value function induced by the contract cannot dip below the horizontal axis. That is, it must lie everywhere above 0. However, by randomizing judiciously (though randomization may not be necessary), the uninformed Bob can always secure a payoff no less than the minimum of the value function. He is not screened out.

By allowing the informed Bob to learn, we make it so that although the value function does dip below the horizontal axis, the value function evaluated at any posterior that may result from learning is positive. The uninformed Bob is screened out.

2.2 Two Examples

Two State, Single Experiment Example

Suppose there are two states, H and L ; and let $\mu := \mathbb{P}(H)$ denote the informed Bob's prior. Regardless of what prior he has, informed Bob has ac-

cess to only one experiment: he can get a signal $s \in \{h, l\}$ such that $\mathbb{P}(h|H) = \mathbb{P}(l|L) = 3/4$, and the cost of performing the experiment is $c = 50$. Letting $(u, d) = (250, 600)$, if informed Bob has $\mu \in (1/3, 2/3)$, he acquires information, as his payoff from doing so is

$$u - d \cdot [\min\{\mu\mathbb{P}(h|H), (1 - \mu)\mathbb{P}(h|L)\} + \min\{\mu\mathbb{P}(l|H), (1 - \mu)\mathbb{P}(l|L)\}] - c = 50.$$

If he has $\mu \notin (1/3, 2/3)$ his payoff is greater than or equal to $u - d/3 = 50 > 0$. But the uninformed Bob's expected payoff is $u - d/2 = -50 < 0$. Thus, the contract (u, d) screens the uninformed Bob.

Two State, Posterior-Separable Cost Example

Suppose there are just two states, the informed Bob has access to any experiment, and his cost of obtaining information is monotone in the Blackwell order. For convenience, we write his cost of obtaining information as a cost defined on his resulting distribution over posteriors F : the cost of acquiring the experiment that produces F is $\Gamma(F) = \kappa \int_0^1 c(x) dF(x) - \kappa c(\mu)$ for some strictly convex function $c: [0, 1] \rightarrow \mathbb{R} \cup \{\infty\}$, $\kappa \in (0, K]$, and $K \in \mathbb{R}_{++}$.⁷

Observe that a contract induces value function $V(x) = u - d \min\{x, 1 - x\} - \kappa c(x) + \kappa c(\mu)$, which is kinked precisely at $x = \frac{1}{2}$. Accordingly, for all κ and c , and for all μ s such that the informed Bob acquires information, the informed Bob's optimal learning will not have support in some interval (x_L, x_H) , where $x_L < \mu < x_H$. Moreover, if $u - \frac{1}{2}d = 0$, $\mathbb{E}_{F_\mu^*} V > 0$ for all $\mu \in [0, 1]$, where F_μ^* is an optimally acquired distribution.⁸ Accordingly,

⁷It is important to note that we are using the fact that C depends not only on the experiment χ , but the prior, μ , as well or else this would be impossible, as pointed out by [Denti et al. \(2022a\)](#) and [Denti et al. \(2022b\)](#) (see also [Mensch \(2018\)](#)).

⁸Note that F_μ^* may be the degenerate distribution on μ .

there is some $\varepsilon > 0$ such that $u - \varepsilon - \frac{1}{2}d < 0$ but $\max_{F_\mu} \mathbb{E}_{F_\mu} V(x) - \varepsilon > 0$ for all $\mu \in [0, 1]$.

The optimal randomization strategy for the uninformed Bob, should he accept the contract, is to announce each state with equal probability, yielding him a payoff of $u - \varepsilon - \frac{1}{2}d < 0$. Thus, only the informed Bob will deliver a theory.

Figure 1 illustrates a collection of value functions induced by a contract that screens the uninformed Bob, when the informed Bob's cost function is the (expected) reduction in (Shannon) entropy. The dotted red curve is the value function for an informed Bob with prior $\mu = \frac{1}{2}$. The grey curves are value functions for various other μ s. They are pointwise increasing as μ gets further from $\frac{1}{2}$. The dotted orange lines are the concavifying lines corresponding to Bob's optimal learning (for the different priors). The Bob with the highest depicted curve ($\mu = .53$ or $.47$) does not acquire any information. Indeed, for our main result to hold, we only need the informed Bob with prior sufficiently close to the kink $\frac{1}{2}$ to "move away from there" by acquiring information.

2.3 An Alternate Specification

Our formal model is not as true to [Olszewski and Sandroni \(2011\)](#) as it could have been, but was chosen instead to cleanly illustrate our point. Let us now quickly explore a variant that parallels their example more closely in order to assure ourselves that our insights are correct.

Suppose now that there are three states $\theta \in \{rr, rb, bb\}$, which captures a situation in which an urn contains two balls, each of which takes one of two possible colors, red (r) or black (b). The timing of the scenario remains the




Figure 1: The value functions (gross of costs) induced by a contract that screens the uninformed Bob. [Try it yourself!](#)

same but now the contract conditions not on the true state but on Alice's evidence about it, which is the realization of a uniform draw from the urn. Let $x := \mathbb{P}(rr)$ and $y := \mathbb{P}(bb)$, denote Bob's beliefs about the state. For simplicity, we assume Bob's cost on experiments produces the uniformly posterior-separable form (as in the previous subsection)

$$\Gamma(F) = \kappa \int_{\Delta} c(x, y) dF(x, y) - \kappa c(\mu_x, \mu_y).$$

Bob has two possible announcements (you will not draw) *red* and *black*. Now, a contract induces value function $V(x, y) = u - \frac{d}{2} \min\{1 + x - y, 1 + y - x\} - \kappa c(x, y) + \kappa c(\mu_x, \mu_y)$. This is kinked at the line $y = x$. Thus, no matter the (finite) κ or c , the informed Bob's optimal learning will have no support within some open superset of $\{(x, y) \in \Delta: y = x\}$. Moreover, if $u - \frac{1}{2}d = 0$, $\mathbb{E}_{F_{\mu}^*} V > 0$ for all $\mu \in \Delta$, where F_{μ}^* is an optimally acquired distribution.

Accordingly, there is some $\varepsilon > 0$ such that $u - \varepsilon - \frac{1}{2}d < 0$ but

$$\max_{F_\mu} \mathbb{E}_{F_\mu} V - \varepsilon > 0$$

for all $\mu \in \Delta$. The uninformed Bob, should he accept the contract, again optimally announces each state with equal probability, yielding him a payoff of $u - \varepsilon - \frac{1}{2}d < 0$. Thus, only the informed Bob will deliver a theory.

2.4 When Uninformed Bob is an SEU-Maximizer

Olszewski and Sandroni (2011) aptly model a scientific theory as knowledge of the state of the world (see our §2.3), which we simplify to knowledge of the distribution of the state of the world (our main specification). Alas, they show that falsifiability is unable to distinguish correct theories from incorrect ones. We modified their model so that the informed expert is not perfectly informed about the state, but instead can acquire information.

This seems reasonable to us, especially if one thinks of science as an incremental process, where hypotheses are refined over time, converging to consensus about the state only in the long-run limit. This raises the question; however, as to how we should think about the uninformed Bob. Perhaps he should be an expected-utility maximizer, himself, with a well-defined subjective prior. Given this, the value of the falsifiability paradigm is how it potentially enables the production of scientific theories only by the Bob who can learn, ensuring progress and eventual knowledge.

Formally, we can adapt Theorem 2.1 to the case where the uninformed Bob is a subjective expected-utility maximizer with a known prior $\rho \in \text{int } \Delta$. Define a **generalized contract** to be $n + 1$ scalars $u > 0$ and $d_i > 0$ ($i = 1, \dots, n$). As before, if Bob accepts the contract he obtains the payment, u , up

front; then announces a state that will not realize. If the announced state θ_i coincides with the realized state, Alice levies the penalty, d_i , on Bob.

We adapt our earlier joint assumption on C and \mathcal{E} by replacing $\frac{1}{n}$ with ρ :

(A') There exist $\varepsilon, \eta > 0$ such that for all $\mu \in B(\rho, \eta)$, there exists $E_\mu \in \mathcal{E}$ that is ε -valuable for μ , and $C(E_\mu, \mu) \leq T$ for some $T \in \mathbb{R}_+$.

Then,

Proposition 2.3. *There exists a generalized contract that screens the uninformed Bob.*

Proof. It suffices to show that there is a solution to

$$-\rho_1 d_1 = -\rho_2 d_2 = \dots = -(1 - \rho_1 - \dots - \rho_{n-1}) d_n,$$

with one degree of freedom. The rest follows the proof of Theorem 2.1, *mutatis mutandis*. Such a solution is

$$d_i = \frac{1 - \rho_1 - \dots - \rho_{n-1}}{\rho_i} d_n \text{ for all } i = 1, \dots, n-1,$$

where d_n is a free variable. The d_i s are well-defined because $\rho \in \text{int } \Delta$. ■

If d_i must equal d for all i —i.e., generalized contracts are forbidden—unless $\rho = \frac{1}{n}$, it may not be possible to screen the uninformed Bob.

2.4.1 Unknown Prior for Uninformed Bob

If the uninformed Bob's prior is also unknown to Alice, there is no contract that is such that informed Bob accepts the contract and uninformed Bob does not, no matter their priors.⁹ This is because for an SEU-maximizing

⁹This statement is true even if we understand “contract” to mean an arbitrary finite menu of state-contingent transfers.

uninformed Bob, rejecting the contract at all priors is equivalent to his outside option payoff 0 being strictly dominant. Thus, it will be for the informed Bob as well.

Nevertheless, a different sort of result is true. We say that informed Bob has a **Rich Learning Set** if for all $\mu \in \text{int}\Delta$, and all $E \in \mathcal{E}$ that produce distributions over posteriors with support on $\text{int}\Delta$, $C(E, \mu) < \infty$.¹⁰ A contract ξ -**screens** the uninformed Bob if the informed Bob prefers to accept it, but the set of priors at which uninformed Bob strictly prefers to reject has Lebesgue measure of at least $1 - \xi$.

Remark 2.4. For all $\xi > 0$, if informed Bob has a rich learning set, there is a contract that ξ -screens the uninformed Bob.

Proof. For any $\xi > 0$, u and d can be chosen such that the induced value function, $V(\mathbf{x}, i)$ is strictly negative on a subset of Δ with Lebesgue measure $1 - \xi$; and ii) is such that $V(\mathbf{x})$ is arbitrarily large for all $\mathbf{x} \in \partial\Delta$. We are done: at most measure ξ of the uninformed Bobs will accept, but all of the informed ones will. ■

Note that such a contract induces the informed Bob to acquire arbitrarily precise information.

References

Andrew Caplin, Mark Dean, and John Leahy. Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political*

¹⁰This assumption is satisfied when the informed Bob's choice of experiments is unrestricted, and the cost is posterior-separable or experimental (under the standard assumptions for such costs).

- Economy*, 130(6):1676–1715, 2022.
- Aubrey Clark and Giovanni Reggiani. Contracts for acquiring information. *arXiv: 2103.03911*, 2021.
- A Philip Dawid. The well-calibrated bayesian. *Journal of the American Statistical Association*, 77(379):605–610, 1982.
- A Philip Dawid. Comment: The impossibility of inductive inference. *Journal of the American Statistical Association*, 80(390):340–341, 1985.
- Eddie Dekel and Yossi Feinberg. Non-bayesian testing of a stochastic prediction. *The Review of Economic Studies*, 73(4):893–906, 2006.
- Tommaso Denti, Massimo Marinacci, and Aldo Rustichini. Experimental cost of information. *American Economic Review*, 112(9):3106–23, 2022a.
- Tommaso Denti, Massimo Marinacci, and Aldo Rustichini. The experimental order on random posteriors. *Mimeo*, 2022b.
- Dean P Foster. A proof of calibration via blackwell’s approachability theorem. *Games and Economic Behavior*, 29(1-2):73–78, 1999.
- Dean P Foster and Rakesh V Vohra. Asymptotic calibration. *Biometrika*, 85(2):379–390, 1998.
- Drew Fudenberg and David K Levine. An easier way to calibrate. *Games and economic behavior*, 29(1-2):131–137, 1999.
- Sergiu Hart. Calibrated forecasts: The minimax proof. *Mimeo*, 2022.
- Benjamin Hébert and Michael Woodford. Neighborhood-based information costs. *American Economic Review*, 111(10):3225–3255, 2021.

- Tai Wei Hu and Eran Shmaya. Expressible inspections. *Theoretical Economics*, 8(2):263–280, 2013.
- Ehud Lehrer. Any inspection is manipulable. *Econometrica*, 69(5):1333–1347, 2001.
- Elliot Lipnowski and Doron Ravid. Predicting choice from information costs. *Mimeo*, 2023.
- 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.
- Jeffrey Mensch. Cardinal representations of information. *Mimeo*, 2018.
- Michèle Müller-Itten, Roc Armenter, and Roc Stangebye. Rational inattention via ignorance equivalence. *Mimeo*, 2021.
- Wojciech Olszewski. Calibration and expert testing. In *Handbook of Game Theory with Economic Applications*, volume 4, pages 949–984. Elsevier, 2015.
- Wojciech Olszewski and Alvaro Sandroni. Contracts and uncertainty. *Theoretical Economics*, 2(1):1–13, 2007.
- Wojciech Olszewski and Alvaro Sandroni. Manipulability of future-independent tests. *Econometrica*, 76(6):1437–1466, 2008.
- Wojciech Olszewski and Alvaro Sandroni. Manipulability of comparative tests. *Proceedings of the National Academy of Sciences*, 106(13):5029–5034, 2009.

- Wojciech Olszewski and Alvaro Sandroni. Falsifiability. *American Economic Review*, 101(2):788–818, April 2011.
- Luciano Pomatto, Philipp Strack, and Omer Tamuz. The cost of information. *Mimeo*, 2020.
- Daniel Rappoport and Valentin Somma. Incentivizing Information Design. *Available at SSRN 3001416*, 2017.
- Alvaro Sandroni, Rann Smorodinsky, and Rakesh V Vohra. Calibration with many checking rules. *Mathematics of operations Research*, 28(1):141–153, 2003.
- Salil Sharma, Elias Tsakas, and Mark Voorneveld. Procuring unverifiable information. *Mimeo*, 2020.
- Eran Shmaya. Many inspections are manipulable. *Theoretical Economics*, 3(3):367–382, 2008.
- Christopher A Sims. Stickiness. In *Carnegie-rochester conference series on public policy*, volume 49, pages 317–356, 1998.
- Christopher A Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003.
- Mark Whitmeyer and Kun Zhang. Buying opinions. *Mimeo*, 2022.
- Nathan Yoder. Designing incentives for heterogeneous researchers. *Journal of Political Economy*, 130(8):2018—2054, 2022.
- Luis Zermeno. A principal-expert model and the value of menus. *Mimeo*, 2011.