

Can One Hear the Shape of a Decision Problem?

Mark Whitmeyer¹

We explore the connection between an agent's decision problem and her ranking of information structures. We find that a finite amount of ordinal data on the agent's ranking of experiments is enough to identify her (finite) set of undominated actions (up to relabeling and duplication) and the beliefs rendering each such action optimal. An additional smattering of cardinal data, comparing the relative value to the agent of finitely many pairs of experiments, identifies her utility function up to an action-independent payoff.

Keywords: Expected Utility, Value of Information, Rationalizability

JEL Classifications: C72; D81; D82; D83

1 Motivation

To what extent does an agent's ranking of information structures specify the decision problem she faces? We know she is a Bayesian expected-utility maximizer and that her conception of the world has finitely many states. We, furthermore, know her prior (though this is inessential) and that only finitely many of her available actions are ever optimal; *viz.*, are undominated. What we do not know is her utility function itself.

¹ Arizona State University. Email: mark.whitmeyer@gmail.com. I thank Wenhao Li, Joseph Whitmeyer, Tom Wiseman, and Kun Zhang for their comments.

The goal of this work is to understand the connection between an agent's value for information and her decision problem. Obviously, knowing an agent's decision problem and prior fully specifies her value for information: her (expected) utility gain from observing any Blackwell experiment (Blackwell (1951), Blackwell (1953)) can be computed with ease. Any two experiments can, therefore, be ranked. Here, we study the inverse problem: suppose we know a finite collection of the agent's comparative ranks of experiments. Does that pin down the agent's utility function?

The three main results of this paper are two positive ones, with a negative one sandwiched in between. In the first, Lemma 3.2, we discover that a finite collection of ranked experiments identifies the agent's undominated actions (up to relabeling and duplication) as well as the regions of beliefs on which each such action is optimal. The second result, Proposition 3.5, reveals an important indeterminacy about the agent's utility function that cannot be resolved with only ordinal rankings of experiments. The third result, Theorem 3.7, states that an additional finite collection of utility differences between experiments (the difference in utils from observing one experiment versus another) identifies the agent's utility function up to a decision-irrelevant payoff.

These results show that an agent's preferences for information reveal a significant amount about her decision problem. A finite amount of ordinal information alone is enough to completely characterize the agent's behavior from observing any information structure. An additional finite number of utility differences yields the agent's exact value for information, for any comparison of information structures. The rationale behind these results is that expected utility maximization imposes a significant amount of structure on how an agent values information.

As two-player (simultaneous-move) games, fixing the behavior of one player, are decision problems from the perspective of the other player, we also find that a finite collection of statements about the intersections of best-response sets is enough to characterize a two-player game, up to inessential details. That is, such a finite collection identifies each player's set of rationalizable actions (up to duplication and relabeling), as well as the sets of beliefs rationalizing each such action. This, in turn, pins down the Nash equilibria of the game.

2 The Question(s) at Hand

There is a decision-maker (DM) encountering a static decision problem, which we define to be a triple (Θ, A, u) consisting of a finite set of possible states Θ , a nonempty set of actions A , and a bounded utility function $u: A \times \Theta \rightarrow \mathbb{R}$. We, the observers, know Θ and the DM's full support prior $\mu \in \Delta(\Theta) \equiv \Delta$. We also know that the DM is a Bayesian expected-utility maximizer whose utility function is bounded, only finitely many of her actions are undominated, and no two undominated actions have identical payoffs in every state.^[2]

Furthermore, we also have a finite collection of the DM's rankings of signals. That is, there is a finite collection of Blackwell experiments π_1, \dots, π_m and rankings of the form $\pi_i \succ \pi_j$ and $\pi_k \sim \pi_l$ ($i, j, k, l \in \{1, \dots, m\}$); meaning the DM strictly prefers experiment π_i to π_j and is indifferent between experiments π_k and π_l , respectively.

Identifying the decision problem: is this data—the finite collection of ranked experiments, the prior, that only finitely many actions are undominated, etc.—enough to pin down the decision problem up to an action-irrelevant payoff (and to relabeling of the actions)?

If this data is insufficient, will a finite supplement of cardinal information suffice in the following sense? For a Blackwell experiment π and a prior μ , let X be the random posterior produced by Bayesian updating, which is distributed according to $F = B(\mu, \pi)$ (where B is the Bayes map).^[3] Then, let $W(\pi) \in \mathbb{R}$ denote the DM's Expected payoff in her decision problem from first observing the outcome of experiment π then taking an optimal action:

$$W(\pi) = \mathbb{E}_F \max_{a \in A} \mathbb{E}_x u(a, \theta).$$

Suppose we also have a finite collection of Blackwell experiments π_1, \dots, π_m , real numbers ζ_1, \dots, ζ_k , and equalities $W(\pi_i) = W(\pi_j) + \zeta_l$. Note that we do not know W , this data consists of statements “the DM's utility gain from π_j to π_i is ζ_l utils.” Now can we identify

^[2] Alternatively, we resign ourselves to being unable to distinguish between actions that are exact duplicates. This is a genericity requirement.

^[3] We borrow this notation from [Denti, Marinacci, and Rustichini \(2022a\)](#) and [Denti, Marinacci, and Rustichini \(2022b\)](#).

the state-dependent utility up to a decision-irrelevant random payoff (and to relabeling of the actions)?

2.1 Recasting the Question

Given the set of states Θ , a set of actions A , and a bounded state-dependent utility function $u: A \times \Theta \rightarrow \mathbb{R}$, the DM's value function in belief $x \in \Delta$ is

$$V(x) := \max_{a \in A} \mathbb{E}_x u(a, \theta).$$

There are $n \geq 2$ states ($|\Theta| = n$), so $\Delta \subset \mathbb{R}^{n-1}$.

An action $a \in A$ is undominated if there exists a belief $x \in \Delta$ such that

$$\mathbb{E}_x u(a, \theta) > \mathbb{E}_x u(b, \theta), \text{ for all } b \in A \setminus \{a\}.$$

As only finitely many actions are undominated, V is the maximum of finitely many hyperplanes, i.e., is convex and piecewise affine.

As noted before, via Bayes' law, given prior $\mu \in \Delta$ any experiment π corresponds to a random vector X supported on Δ whose expectation is the DM's prior μ . Accordingly, we have a finite collection ($m \in \mathbb{Z}_{++}$) of equalities and strict inequalities of the form $\mathbb{E}_{F_i} V(X_i) > \mathbb{E}_{F_j} V(X_j)$ and $\mathbb{E}_{F_k} V(X_k) = \mathbb{E}_{F_l} V(X_l)$ ($i, j, k, l \in \{1, \dots, m\}$). We call this a collection of ordered expectations.

Identifying the decision problem: is this data—the collection of ordered expectations, and that V is the maximum of finitely many affine functions—enough to pin down the decision problem? That is, is there such a collection that is satisfied by a piecewise-affine V ; and, moreover, if W satisfies the collection, then $W = V + \varphi$, where $\varphi: \Delta \rightarrow \mathbb{R}$ is affine? What if we also have a finite supplement of additional *cardinal* data of the form $\mathbb{E}_{F_i} V(X_i) = \mathbb{E}_{F_j} V(X_j) + \zeta_l$? Will that pin down $V + \varphi$?

3 Answering the Question

Any such value function V can be projected onto the belief simplex, Δ , which yields a finite collection C of polytopes of full dimension C_i ($i = 1, \dots, t$), where t is the number of undominated actions in A . Formally, for each undominated $a_i \in A$ ($i = \{1, \dots, t\}$),

$$C_i := \{x \in \Delta \mid \mathbb{E}_x u(a_i, \theta) = V(x)\}.$$

Action a_i is optimal for any belief $x \in C_i$ and uniquely optimal for any belief in the relative interior of C_i (for any $x \in \text{int } C_i$). The collection C is a regular polyhedral subdivision (henceforth, just “subdivision”) of Δ . We say that function V induces subdivision C .

Definition 3.1. *A collection of ordered expectations identifies subdivision C if there exists a V that satisfies the collection that induces C , and any V that satisfies the collection of ordered expectations induces C .*

We can always find a subdivision-identifying collection of ordered expectations.

Lemma 3.2. *For any C , there exists a collection of ordered expectations that identifies it.*

Proof. Fix an arbitrary C . Observe that for any element $C_i \in C$, there exists a random vector X_i with finitely supported distribution F_i such that part of the support of F_i is the extreme points of C_i , i.e., $\text{extreme } C_i \subseteq \text{supp } F_i$. Take another \hat{X}_i , with distribution \hat{F}_i , where \hat{F}_i is a mean-preserving contraction of F_i that is identical to F_i except that the portion of the measure placed on the extreme points of C_i is collapsed to its barycenter, which by construction is in $\text{int } C_i$. Then, $\mathbb{E}_{F_i} V(X_i) = \mathbb{E}_{\hat{F}_i} V(\hat{X}_i)$ if and only if V is affine on C_i . Accordingly, for a C with t elements, t such equalities implies V is affine on each $C_i \in C$.

Next, for any distinct $C_i, C_j \in C$ with $\dim(C_i \cap C_j) = n - 2$ and any $x' \in \text{int}(C_i \cap C_j)$, there exists an X_k with finitely supported distribution F_k such that $x' \in \text{supp } F_k$. Take another \hat{X}_k , with distribution \hat{F}_k , where \hat{F}_k is a mean-preserving spread of F_k that is identical to F_k except that the portion of the measure placed on x' is split into two points $x_1 \in \text{int } C_i$ and $x_2 \in \text{int } C_j$. Then, $\mathbb{E}_{F_k} V(X_k) < \mathbb{E}_{\hat{F}_k} V(\hat{X}_k)$ if and only if V is not affine on $C_i \cup C_j$. As there are only finitely many pairings of distinct C_i s, finitely many such strict inequalities imply that V is not affine on any such pair.

As V is convex, any V that satisfies the stated equalities and strict inequalities must induce C . Moreover, for any C , there is a V that induces it (indeed this is the definition of a regular polyhedral subdivision). ■

Figures 1 and 2 illustrate the proof of Lemma 3.2. In Figure 1, revealing that the DM (with prior on the empty-interior green point) is indifferent between the distribution over posteriors with support on the black dots and the distribution over posteriors with support on the black point in the blue region and the purple x identifies that V is affine




Figure 1: Showing V is affine on the red region

on the red region. In Figure 2, revealing that the DM (with prior on the empty-interior green point) strictly prefers the distribution over posteriors with support on the black and purple x s to the distribution over posteriors with support on the black dot and the purple x identifies that V is not affine on the red and blue regions.

Next we uncover what the equivalence of value functions means in terms of the state-dependent utilities. In the next lemma, let value function \hat{V} correspond to utility \hat{u} .

Lemma 3.3. $\hat{V}(x) = V(x) + \varphi(x)$ for affine φ if and only if the sets of undominated actions for both $\hat{u}(a, \theta)$ and $u(a, \theta)$ are the same up to relabeling the actions; and $\hat{u}(a, \theta) = u(a, \theta) + \gamma(\theta)$ for some $\gamma: \Theta \rightarrow \mathbb{R}$ for all undominated actions up to relabeling the actions.

Proof. (\Rightarrow) If $\hat{u}(a, \theta) = u(a, \theta) + \gamma(\theta)$, then

$$\hat{V}(x) = \max_{a \in A} \mathbb{E}_x \hat{u}(a, \theta) = \max_{a \in A} \mathbb{E}_x [u(a, \theta) + \gamma(\theta)] = V(x) + \mathbb{E}_x \gamma(\theta) = V(x) + \varphi(x).$$

(\Leftarrow) Suppose for the sake of contraposition that there is no γ such that $\hat{u}(a, \theta) = u(a, \theta) + \gamma(\theta)$ for all undominated actions. If the subdivisions induced by V and \hat{V} are not the same, we are done (see, e.g., Lemma 2.1 in [Whitmeyer \(2023\)](#)). Now suppose they are the same,




Figure 2: Showing V is not affine on the red and blue regions

with the same actions corresponding to each element (by relabeling). By assumption, there exist two undominated actions a_1 and a_2 for which $\hat{u}(a_1, \theta) = u(a_1, \theta) + \gamma_{a_1}(\theta)$ and $\hat{u}(a_2, \theta) = u(a_2, \theta) + \gamma_{a_2}(\theta)$, where $\gamma_{a_1} \neq \gamma_{a_2}$. Thus,

$$\mathbb{E}_x \gamma_{a_1}(\theta) = \varphi_1(x) \neq \varphi_2(x) = \mathbb{E}_x \gamma_{a_2}(\theta),$$

so $\hat{V}(x) = V(x) + \varphi_1(x)$ for all $x \in C_1$ and $\hat{V}(x) = V(x) + \varphi_2(x)$ for all $x \in C_2$. ■

Now that we know we can find a collection of ordered expectations that identifies C we turn our attention to identifying V . The crucial part of the proof of the next result is the observation that C pins down V (plus affine φ) up to one degree of freedom, but that this last remaining indeterminacy cannot be eliminated by a collection of ordered expectations.

Definition 3.4. *A collection of ordered expectations identifies value function V , if there exists a value function that satisfies the collection, and any W that satisfies the collection equals V plus some affine function φ .*

Proposition 3.5. *For any non-affine V , no collection of ordered expectations identifies it.*

Proof. Fix an arbitrary V . From Lemma 3.2, there exists a collection of ordered expectations that identifies the C it induces. On any $C_i \in C$, $V(x) = \alpha_i \cdot x + \beta_i$. It is immediate that if V satisfies a collection of ordered expectations and induces C , then so too does any ψV ($\psi > 0$). As V is not affine, we are done. ■

In some sense this result is to be expected. How could purely ordinal data identify V ? Nevertheless, a few pieces of cardinal information will do. We term an equality of the form $\mathbb{E}_{F_i} V(X_i) = \mathbb{E}_{F_j} V(X_j) + \zeta_l$ a utility difference.

Definition 3.6. *A collection of ordered expectations and utility differences identifies value function V , if there exists a value function that satisfies the collection, and any value function that satisfies the collection equals V plus some affine function φ .*

Theorem 3.7. *For any V , there exists a collection of ordered expectations and utility differences that identifies it.*

Proof. Fix an arbitrary V and its induced C . Take any distinct $C_i, C_j \in C$ with $\dim(C_i \cap C_j) = n - 2$. WLOG we may specify $V = 0$ for all $x \in C_i$, as our identification is up to the addition of an affine function. By construction, $V(x) = \psi(\alpha \cdot x + \beta)$ ($\alpha \in \mathbb{R}^{n-1}$, $\beta \in \mathbb{R}$, $\psi \in \mathbb{R}_{++}$) for all $x \in C_j$, with $\alpha \cdot x + \beta \geq 0$ for all $x \in C_j$. Moreover, $\alpha \cdot x + \beta > 0$ for all $x \in C_j \setminus C_i$.

Take two points $x_i \in \text{int } C_i$ and $x_j \in \text{int } C_j$ such that the set $\lambda x_i + (1 - \lambda) x_j \in C_i \cup C_j$ for any $\lambda \in [0, 1]$. Let $\hat{x} = \text{int conv } \{x_i, x_j\} \cap \text{int } C_i$. WLOG $\mu \in \text{int conv } \{\hat{x}, x_j\}$. Then, letting F be the binary distribution with support $\{x_i, x_j\}$ and \hat{F} be the binary distribution with support $\{\hat{x}, x_j\}$, $\mathbb{E}_F V(X) > \mathbb{E}_{\hat{F}} V(\hat{X})$. Define $\zeta = \mathbb{E}_F V(X) - \mathbb{E}_{\hat{F}} V(\hat{X}) > 0$.

Finally, observe that, letting p and q be defined implicitly as

$$px_j + (1 - p)x_i = \mu = qx_j + (1 - q)\hat{x},$$

the equation (in ψ)

$$p\psi(\alpha \cdot x_j + \beta) = q\psi(\alpha \cdot x_j + \beta) + \zeta$$

has a unique solution

$$\psi = \frac{\zeta}{(p - q)(\alpha \cdot x_j + \beta)} > 0.$$

Consequently, we have identified V on C_j via this single equality. As there are only finitely many pairs of C_i and C_j , we are done. \blacksquare

As identification of V yields the value of information (in utils), $\mathbb{E}_F V - V(\mu)$,

Corollary 3.8. *There exists a collection of ordered expectations and utility differences that identifies the value of information for the DM.*

4 The Irrelevance of the Prior

Until now we have assumed that we know the DM's prior. What if we do not? Take any (full-support) prior $\mu \in \text{int } \Delta$ and belief $x \in \Delta$, that corresponds to the realization, s of a statistical experiment. It is well-known that the ratios of conditional probabilities of that s (conditioned on the states) corresponding to that belief x are well-defined functions of the prior. Consequently, for any prior $\mu \in \text{int } \Delta$ and any decision problem whose value function induces C , the extreme points of any $C_i \in C$ correspond uniquely to the ratios of conditional probabilities of a number of columns (one column per extreme point of C_i) in a statistical experiment.

To put differently, for any $\mu \in \text{int } \Delta$, value function V that induces C , and any $C_i \in C$, the set of extreme points of C_i corresponds uniquely to a family of sub-stochastic matrices, \tilde{C}_i , defined as follows. Denoting by $\underline{\Delta}(S)$ the set of sub-probability measures on a set of signal realizations S , and x_s the DM's posterior after observing signal realization s ,

$$\tilde{C}_i := \{ \tilde{\pi} : \Theta \rightarrow \underline{\Delta}(S_i) : x_s \in \text{extreme}(C_i) \forall s \in S_i, \cup_{s \in S_i} \{x_s\} = \text{extreme}(C_i) \}.$$

By construction, for a decision problem whose value function induces subdivision C , at prior $\mu \in \text{int } \Delta$, the DM is indifferent between any experiment π possessing $\tilde{\pi} \in \tilde{C}_i$ as a submatrix and any experiment $\tilde{\pi}'$ that is the same as π except that $\tilde{\pi}$ is replaced by an arbitrary garbling. We term \tilde{C}_i a spectral element, and the set $\tilde{C} := \{ \tilde{C}_1, \dots, \tilde{C}_t \}$ the spectral subdivision. Importantly, a spectral subdivision is prior-free: it is a set of families of matrices.

We call a finite collection of Blackwell experiments π_1, \dots, π_m and rankings of the form $\pi_i \succ \pi_j$ and $\pi_k \sim \pi_l$ ($i, j, k, l \in \{1, \dots, m\}$) a collection of ranked experiments.

Definition 4.1. *A collection of ranked experiments identifies spectral subdivision \tilde{C} if*

for any prior $\mu \in \text{int } \Delta$ there exists a decision problem that satisfies the collection that induces \tilde{C} , and any decision problem and prior that satisfies the collection induces \tilde{C} .

We immediately deduce a corollary of Lemma 3.2.

Corollary 4.2. *For any spectral subdivision \tilde{C} , there exists a collection of ranked experiments that identifies it.*

We can also add some cardinal information. Recall that for a Blackwell experiment π , $W(\pi) \in \mathbb{R}$ denotes the DM's Expected payoff in her decision problem from first observing the outcome of experiment π then taking an optimal action. We term an equality of the form $W(\pi_i) = W(\pi_j) + \zeta_i$ a utility difference. Recall from Corollary 3.8 that if we knew μ , we could identify the DM's value for information. It turns out, we do not even need to know μ to price (in utils) information for the agent. That is, the precise utility gain from *any* experiment is identified.

We normalize the value of no information to 0 so that the value of information of experiment π is merely $W(\pi)$. Then,

Definition 4.3. *A collection of ranked experiments and utility differences identifies the value of information W if for any prior $\mu \in \text{int } \Delta$ there exists a decision problem such that the value of an experiment π is $W(\pi)$, and any decision problem and prior satisfying the collection has value of information W .*

Theorem 4.4. *There exists a collection of ranked experiments and utility differences that identifies the value of information for the agent.*

Proof. Take an arbitrary prior $\mu' \in \text{int } \Delta$ and a decision problem that induces value function V . By Theorem 3.7 there exists a collection of ordered expectations and utility differences that identifies it and, hence, W for this DM. Moreover, take the collections specified in the proofs of Proposition 3.2 and Theorem 3.7, and observe that the ranked experiments corresponding to the ordered expectations and the utility differences identify value function V^μ for each $\mu \in \text{int } \Delta$. Moreover, by construction, for all such priors μ , the value of information in the corresponding decision problem (with value function V^μ) is the same and, thus, equal to W .

Accordingly, it suffices to show that for an arbitrary prior $\mu' \in \text{int } \Delta$, there exists a collection of ranked experiments and utility differences that identifies value of information for the agent. This is Corollary 3.8, so we are done. ■

5 Identifying a Two-Player Game

There is a tight connection between two-player (simultaneous move) games and decision problems. In such games, the other player takes on the role of nature, generating any uncertainty by randomizing over his strategies. Given this, it is easy to construct analogs of the earlier identification results.

Formally, a finite two-player (normal form) game is $G = (\{1, 2\}, A, u)$ consisting of two players (1 and 2), a finite set of (pure) action profiles $A = A_1 \times A_2$ and a profile of utilities for the players $u = (u_1, u_2)$, with $u_i: A \rightarrow \mathbb{R}$ ($i = 1, 2$). $\Sigma = \Sigma_1 \times \Sigma_2$ denotes the set of (mixed) strategy profiles, with $\Sigma_i = \Delta A_i$ ($i = 1, 2$). As above, we stipulate that there are no duplicate actions for either player. As before, this is an imposition of genericity.

We define player 1's best response correspondence, in $\sigma_2 \in \Sigma_2$, $b_1(\sigma_2)$, in the standard manner:

$$b_1(\sigma_2) := \left\{ a_1 \in A_1 : \mathbb{E}_{\sigma_2} u(a_1, a_2) = \max_{a'_1 \in A_1} \mathbb{E}_{\sigma_2} u(a'_1, a_2) \right\}.$$

We define player 2's best response correspondence $b_2(\sigma_1)$ analogously. We term a finite collection of the form $b_1(\sigma_2^i) \cap b_1(\sigma_2^j) \neq \emptyset$, $b_2(\sigma_1^i) \cap b_2(\sigma_1^j) \neq \emptyset$, $b_1(\sigma_2^k) \cap b_1(\sigma_2^l) = \emptyset$, and $b_2(\sigma_1^k) \cap b_2(\sigma_1^l) = \emptyset$ a collection of best-response comparisons.

As in the decision problem above, Σ_2 is subdivided into a finite collection C_1 of polytopes of full dimension C_1^i , where for each undominated $a_1 \in A_1$,

$$C_1^i := \left\{ \sigma_2 \in \Sigma_2 \mid a_1^i \in b_1(\sigma_2) \right\}.$$

Σ_1 is likewise subdivided into collection C_2 . We call C_i player i 's subdivision and say that the game induces subdivision pair (C_1, C_2) if the game is such that player 1's subdivision is C_1 and player 2's subdivision is C_2 .

Definition 5.1. *A collection of best-response comparisons identifies game G if there exists a game whose best-response correspondences satisfy the collection, and any game whose best-response correspondences satisfy the collection induces the same subdivision pair (C_1, C_2) .*

Mirroring the argument for Lemma 3.2, we have

Theorem 5.2. *For any game there exists a collection of best-response comparisons that identifies it.*

Proof. As the proof is virtually identical to that for Lemma 3.2, we merely sketch it.⁴ Take an arbitrary element $C_1^i \in C_1$. It is a polytope with a finite set of extreme points (its vertices) $\{\sigma_2^1, \sigma_2^2, \dots, \sigma_2^k\}$. Accordingly, for some $\sigma_2^* \in \text{int } C_1^i$ the collection

$$b(\sigma_2^1) \cap b(\sigma_2^2) \cap \dots \cap b(\sigma_2^k) \cap b(\sigma_2^*) \neq \emptyset$$

implies that there is an action that is a best response to any $\sigma_2 \in C_1^i$. There are only finitely many elements in C_1 , so finitely many more such collections imply that there are actions that are best responses, respectively, to any $\sigma_2 \in C_1^i$ for any $C_1^i \in C_1$. Finally, for any distinct pair C_1^i and C_1^j (of which there are finitely many) a comparison

$$b(\sigma_2^\heartsuit) \cap b(\sigma_2^\diamond) = \emptyset,$$

where $\sigma_2^\heartsuit \in \text{int } C_1^i$ and $\sigma_2^\diamond \in \text{int } C_1^j$, distinguishes between the respective best responses. ■

6 Related Work

This paper belongs among those studying the value of information in decision problems. That literature includes Ramsey (1990), Athey and Levin (2018), and Azrieli and Lehrer (2008); as well as Radner and Stiglitz (1984), De Lara and Gilotte (2007), Chade and Schlee (2002), and Whitmeyer (2024), which four all study the marginal valuation of information. We also make use of the connection between decision problems and polyhedral subdivisions. Green and Osband (1991) reveal this connection, and Lambert (2019) and Frongillo and Kash (2021) make use of it when studying elicibility of various forms of information.

A number of other works seek to understand what information a DM obtains based on her choice behavior in a (known) decision problem. Dillenberger, Lleras, Sadowski, and Takeoka (2014) and Lu (2019) are two such papers. De Oliveira, Denti, Mihm, and Ozbek (2017) and Caplin and Dean (2015) are two others, both of which endogenize the DM's information structure. Libgober (2021) shows how to identify information structures from a population's beliefs.

Finally, this work is fundamentally a revealed preference exercise (Samuelson (1938) and Afriat (1967)): given a choice paradigm what can a set of observed choices by an agent tell you about the agent (and the model)? More recently, Caplin and Martin (2021)

⁴ In fact it is even easier, as we are no longer constrained by Bayes' law.

and [Caplin, Martin, and Marx \(2023\)](#) ask what choice data reveals about the information a DM could have obtained and [Caplin and Dean \(2015\)](#) and [Chambers, Liu, and Rehbeck \(2020\)](#) whether a DM's behavior is consistent with costly information acquisition. This paper, instead, asks what preferences over information structures reveal about a DM's choice problem.

References

- Sydney N Afriat. The construction of utility functions from expenditure data. *International economic review*, 8(1):67–77, 1967.
- Susan Athey and Jonathan Levin. The value of information in monotone decision problems. *Research in Economics*, 72(1):101–116, 2018.
- Yaron Azrieli and Ehud Lehrer. The value of a stochastic information structure. *Games and Economic Behavior*, 63(2):679–693, 2008.
- David Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102, Berkeley, Calif., 1951. University of California Press.
- David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953.
- 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. Comparison of decisions under unknown experiments. *Journal of Political Economy*, 129(11):3185–3205, 2021.
- Andrew Caplin, Daniel J Martin, and Philip Marx. Rationalizable learning. *Mimeo*, 2023.
- Hector Chade and Edward Schlee. Another look at the Radner–Stiglitz nonconcavity in the value of information. *Journal of Economic Theory*, 107(2):421–452, 2002.
- Christopher P Chambers, Ce Liu, and John Rehbeck. Costly information acquisition. *Journal of Economic Theory*, 186:104979, 2020.

- Michel De Lara and Laurent Gilotte. A tight sufficient condition for Radner–Stiglitz nonconcavity in the value of information. *Journal of Economic Theory*, 137(1):696–708, 2007.
- Henrique De Oliveira, Tommaso Denti, Maximilian Mihm, and Kemal Ozbek. Rationally inattentive preferences and hidden information costs. *Theoretical Economics*, 12(2): 621–654, 2017.
- Tommaso Denti, Massimo Marinacci, and Aldo Rustichini. Experimental cost of information. *American Economic Review*, 112(9):3106–23, September 2022a.
- Tommaso Denti, Massimo Marinacci, and Aldo Rustichini. The experimental order on random posteriors, 2022b.
- David Dillenberger, Juan Sebastián Lleras, Philipp Sadowski, and Norio Takeoka. A theory of subjective learning. *Journal of Economic Theory*, 153:287–312, 2014.
- Rafael M Frongillo and Ian A Kash. General truthfulness characterizations via convex analysis. *Games and Economic Behavior*, 130:636–662, 2021.
- Edward J. Green and Kent Osband. A revealed preference theory for expected utility. *The Review of Economic Studies*, 58(4):677–695, 06 1991.
- Nicholas S. Lambert. Elicitation and evaluation of statistical forecasts. *Mimeo*, 06 2019.
- Jonathan Libgober. Hypothetical beliefs identify information. *Mimeo*, 2021.
- Jay Lu. Bayesian identification: a theory for state-dependent utilities. *American Economic Review*, 109(9):3192–3228, 2019.
- Roy Radner and Joseph Stiglitz. A nonconcavity in the value of information. In *M. Boyer, R. Kihlstrom (Eds.), Bayesian Models of Economic Theory*, pages 33–52. Elsevier, Amsterdam, 1984.
- F. P. Ramsey. Weight or the value of knowledge. *The British Journal for the Philosophy of Science*, 41(1):1–4, 1990.
- Paul A Samuelson. A note on the pure theory of consumer’s behaviour. *Economica*, 5 (17):61–71, 1938.

Mark Whitmeyer. Making information more valuable. *Mimeo*, 2023.

Mark Whitmeyer. Call the dentist! a (con-)cavity in the value of information. *Mimeo*, 2024.