

Experimental Cost of Information:

Online Appendix

Tommaso Denti Massimo Marinacci Aldo Rustichini

A Properties of the Bayes map

We review basic properties of the Bayes map $B : \Delta \times \mathcal{E} \rightarrow \Delta^2$ that we use throughout the appendix. Most results are known in statistical decision theory; they can be found in monographs such as Torgersen (1991). For the reader's convenience, here we provide a self-contained presentation. See Denti, Marinacci and Rustichini (2022) for additional results and proofs.

We start by discussing the algebraic properties of the Bayes map. Let $B_P : \Delta \rightarrow \Delta^2$ and $B_\pi : \mathcal{E} \rightarrow \Delta^2$ be its P -section and π -section, respectively.

Lemma 1. *(i) The range of B_π is Δ_π^2 . (ii) The function B_P is injective.*

Property (i) can be decomposed into two parts. First, (i) states that the range of B_π is included by Δ_π^2 . This is a manifestation of the so-called “martingale property” of Bayesian updating: The expected posterior is equal to the prior, i.e., the barycenter of $B(\pi, P)$ is π itself. Second, (i) states that Δ_π^2 is included by the range of B_π : every random posterior is generated by some experiment. This property comes from the richness of the signal space. Overall, (i) and (ii) imply that the map $\mu \mapsto \bar{\mu}$ is the left inverse of B_P .

We turn to the ordinal properties of the Bayes map. Via this map, the Blackwell and convex orders are related in the next result, which extends to priors with full support a classic result for uniform priors.

Lemma 2. *For all $P, Q \in \mathcal{E}$, the following conditions are equivalent:*

- (i) $P \succeq_b Q$;
- (iii) $B(\pi, P) \succeq_{cv} B(\pi, Q)$ for some $\pi \in \Delta_+$;
- (iv) $B(\pi, P) \succeq_{cv} B(\pi, Q)$ for all $\pi \in \Delta_+$.

An implication is that for all $\pi \in \Delta_+$ and $\mu \in \Delta_\pi^2$, the set of experiments $B_\pi^{-1}(\mu)$ is an equivalence class in the Blackwell order. When π does not have full support, we only have that (i) implies (iv).

Next we describe the convexity properties of the Bayes map:

Lemma 3. *For $\alpha \in [0, 1]$, $\pi, \rho \in \Delta$, and $P, Q \in \mathcal{E}$, the following conditions hold:*

$$(i) \alpha B(\pi, P) + (1 - \alpha)B(\rho, P) \succeq_{cv} B(\alpha\pi + (1 - \alpha)\rho, P);$$

$$(ii) \alpha B(\pi, P) + (1 - \alpha)B(\pi, Q) \succeq_{cv} B(\pi, \alpha P + (1 - \alpha)Q).$$

The intuition is that mixing random posteriors corresponds to acquiring information on the basis of a coin toss whose realization the decision maker observes. Mixing experiments corresponds to acquiring information on the basis of a coin toss whose realization the decision maker *does not* observe.

Finally, we report the continuity properties of the Bayes map:

Lemma 4. (i). If $\pi_n \rightarrow \pi$ and $B(\pi^*, P_n) \rightarrow B(\pi^*, P)$, then $B(\pi_n, P_n) \rightarrow B(\pi, P)$.

(ii). If $P_n \rightarrow P$ and $B(\pi, P_n) \rightarrow \mu$, then $\mu \succeq_{cv} B(\pi, P)$.

We conclude with two technical results that complement Lemmas 3 and 4.

Lemma 5. Let $\pi \in \Delta$ and $\mu, \nu \in \Delta_\pi^2$. There are $P, Q \in \mathcal{E}$ such that (i) $B(\pi, P) = \mu$, (ii) $B(\pi, Q) = \nu$, and (iii) $B(\pi, \alpha P + (1 - \alpha)Q) = \alpha\mu + (1 - \alpha)\nu$ for all $\alpha \in [0, 1]$.

Lemma 6. Let (μ_n) be a sequence in Δ_π^2 with $\mu_n \rightarrow \mu$. There is a sequence of experiments (P_n) such that (i) $P_n \rightarrow P$, (ii) $B(\pi, P_n) = \mu_n$ for every n , and (iii) $B(\pi, P) = \mu$.

The two lemmas above use the richness of the signal space.

B Proofs of the results in the main text

Proof of Proposition 2. Fix a cost on experiments $h : \mathcal{E} \rightarrow [0, \infty]$ and a length $n \in \mathbb{N} \cup \{\infty\}$. Assume that $c := c_{h^n}$ satisfies FFI.

Fix a full support prior π^* (for example, the uniform distribution). Let $\hat{\mathcal{E}}^m$ be the set of fully informative sequential experiments of length $m \in \mathbb{N}$:

$$\hat{\mathcal{E}}^m = \left\{ P^m \in \mathcal{E}^m : B(\pi^*, P^m) = \sum_{\theta} \pi^*(\theta) \delta_{\theta}^2 \right\}.$$

By Lemma 2, for all $\pi \in \Delta$,

$$\hat{\mathcal{E}}^m \subseteq \left\{ P^m \in \mathcal{E}^m : B(\pi, P^m) = \sum_{\theta} \pi(\theta) \delta_{\theta}^2 \right\}.$$

Next we derive an implication of FFI for the cost of full information.

Claim 1. For all $\theta \in \Theta$ and $\epsilon > 0$, there are $m \in \mathbb{N}$ and $P^m \in \hat{\mathcal{E}}^m$ such that

$$h(\delta_{\theta}, P^m) \leq \epsilon.$$

In addition, if $n < \infty$, we can choose $m = n$.

Proof of the claim. FFI implies that

$$\lim_{\alpha \rightarrow 1} c \left(\alpha \delta_\theta^2 + (1 - \alpha) \sum_{\tau} \pi^*(\tau) \delta_\tau^2 \right) = 0.$$

Thus, for every $\eta > 0$ there exists $\alpha \in (1/2, 1)$ such that

$$c \left(\alpha \delta_\theta^2 + (1 - \alpha) \sum_{\tau} \pi(\tau) \delta_\tau^2 \right) < \eta.$$

Because $c = c_{h^n}$, we can find $m \in \mathbb{N}$ and $P^m \in \hat{\mathcal{E}}^m$ such that

$$\alpha h(\delta_\theta, P^m) + (1 - \alpha) h(\pi, P^m) \leq \eta.$$

In the case in which $n < \infty$, we can choose $m = n$.

Because $h \geq 0$ and $\alpha \geq 1/2$, we deduce that

$$h(\delta_\theta, P^m) \leq 2\eta.$$

By choosing $\eta = \epsilon/2$, we obtain that $h(\delta_\theta, P^m) \leq \epsilon$. □

Next we describe the main construction of the proof. To simplify the exposition, we order the states: $\Theta = \{\theta_1, \dots, \theta_k\}$ where $k \geq 2$ is the cardinality of Θ .

Fix $\alpha \in (0, 1)$. For every $i = 1, \dots, k - 1$, we invoke Claim 1 to choose $m_i \in \mathbb{N}$ and $P^{m_i} \in \hat{\mathcal{E}}^{m_i}$ such that

$$h(\delta_{\theta_i}, P^{m_i}) \leq \alpha^2. \tag{12}$$

We adopt the convention that P^{m_k} corresponds to not experimenting; thus, in particular, its cost is zero. As in Claim 1, if $n < \infty$, we can choose $m_i = n$.

We inductively construct a sequential experiment Q^m as follows:

- Begin with running P^{m_1} .
- While running P^{m_i} (with $i = 1, \dots, k - 1$), if the posterior probability that the state is θ_i ever falls below α , switch to $P^{m_{i+1}}$.

To clarify the construction, note that while running $P^{m_{k-1}}$, if the posterior probability that the state is θ_{k-1} ever falls below α , then the agent stops experimenting. Moreover, observe that the length of Q^m satisfies the following properties:

- If $n < \infty$, then $m \leq k \cdot n < \infty$.

- If $n = 1$, then $m = 1$.
- If Θ is binary, then $m \leq n$.

Next we characterize the informational content of Q^m .

Claim 2. For every $\pi \in \Delta$, every $\epsilon \in (0, 1/2)$, and every $i \in \{1, \dots, k-1\}$,

$$Q_\pi^m(\{x^m \in X^m : q_{x^m}(\theta_i) \in (\epsilon, 1 - \epsilon)\}) \leq \frac{\alpha}{\epsilon}$$

where $q_{x^m} \in \Delta$ is the posterior belief generated by terminal history x^m , prior π , and sequential experiment Q^m .

Proof of the claim. Define $A^m \subseteq X^m$ by

$$A^m = \{x^m \in X^m : q_{x^m}(\theta_i) \in (\epsilon, 1 - \epsilon)\}.$$

Let $B^l \subseteq X^l$ be the set of histories of length $l = 1, \dots, m-1$ at which the agent makes the decision to abandon P^{m_i} . Let $B^{l,m} \subseteq X^m$ be the set of terminal histories along which the agent makes the decision to abandon P^{m_i} at a history of length $l = 1, \dots, m-1$:

$$B^{l,m} = \{x^m \in X^m : x^l \in B^l\}$$

Let $C^m \subseteq X^m$ be the set of terminal histories along which the agent never makes the decision to abandon P^{m_i} :

$$C^m = \bigcap_{l=1}^{m-1} \{x^m \in X^m : x^m \notin B^{l,m}\}.$$

Every $x^m \in C^m$ is (at least partially) generated by some P^{m_j} (with $j = 1, \dots, i$). Since each P^{m_j} is fully informative, we have that

$$Q_\pi^m(A^m) = Q_\pi^m(A^m \setminus C^m).$$

We observe that

$$\begin{aligned} \int_{X^m \setminus C^m} q_{x^m}(\theta_i) dQ_\pi^m(x^m) &= \sum_{l=1}^{m-1} \int_{B^{l,m}} q_{x^m}(\theta_i) dQ_\pi^m(x^m) \\ &= \sum_{l=1}^{m-1} \int_{B^l} q_{x^l}(\theta_i) dQ_\pi^l(x^l) \leq \alpha \end{aligned}$$

where the second equality follows from the martingale property of Bayesian updating, and

the last inequality is a consequence of $q_{x^l}(\theta_i) \leq \alpha$ (recall the construction of Q^m). Thus,

$$\begin{aligned} \alpha &\geq \int_{X^m \setminus C^m} q_{x^m}(\theta_i) dQ_\pi^m(x^m) \\ &\geq \int_{A^m \setminus C^m} q_{x^m}(\theta_i) dQ_\pi^m(x^m) \geq \epsilon Q_\pi^m(A^m \setminus C^m). \end{aligned}$$

Since $Q_\pi^m(A^m) = Q_\pi^m(A^m \setminus C^m)$, we conclude that $Q_\pi^m(A^m) \leq \alpha/\epsilon$. \square

Next we prove a bound on the cost of Q^m . To do so, we need some additional notation. For every $i \in \{1, \dots, k\}$, we inductively construct a sequential experiment $Q^{(m,i)}$ as follows:

- Begin with running P^{m_i} .
- While running P^{m_j} (with $j = i, \dots, k-1$), if the posterior probability that the state is θ_j ever falls below α , switch to $P^{m_{j+1}}$.

Note that $Q^m = Q^{(m,1)}$, while $Q^{(m,k)}$ corresponds to the choice of not experimenting.

Claim 3. For every $\pi \in \Delta$ and every $i = 1, \dots, k$,

$$h\left(\pi, Q^{(m,i)}\right) \leq 2\alpha(k-i).$$

Proof of the claim. We proceed by induction on i . First observe that trivially,

$$h\left(\pi, Q^{(m,k)}\right) = 0$$

since $Q^{(m,k)}$ corresponds to the choice of not experimenting.

Now suppose that the result is true for $j = i+1$ and all $\pi \in \Delta$. For every $l = 1, \dots, m_i - 1$, define

$$\begin{aligned} A^l &= \left\{x^l \in X^l : p_{x^t}(\theta_i) > \alpha \text{ for all } t = 1, \dots, l\right\} \\ B^l &= \left\{x^l \in X^l : p_{x^l}(\theta_i) \leq \alpha \text{ and } x^{l-1} \in A^{l-1}\right\} \end{aligned}$$

where $p_{x^l} \in \Delta$ is the posterior belief generated by history x^l , prior π , and sequential experiment P^{m_i} ; we adopt the convention that $A^0 = X^0$. In other terms, in the sequential experiment $Q^{(m,i)}$, $A^l \subseteq X^l$ is the set of histories of length $l = 1, \dots, m_i - 1$ along which the agent never makes the decision to abandon P^{m_i} ; $B^l \subseteq X^l$ is the set of histories of length $l = 1, \dots, m_i - 1$ along which the agent makes the decision to abandon P^{m_i} .

Simple algebra shows that

$$h\left(\pi, Q^{(m,i)}\right) = h(P_{x^0}) + \sum_{l=1}^{m_i-1} \left(\int_{A^l} h(P_{x^l}) dP_\pi^l(x^l) + \int_{B^l} h(p_{x^l}, Q^{(m,i+1)}) dP_\pi^l(x^l) \right).$$

By the inductive hypothesis, we have that

$$\sum_{l=1}^{m_i-1} \int_{B^l} h(p_{x^l}, Q^{(m,i+1)}) dP_\pi^l(x^l) \leq 2(k-i-1)\alpha. \quad (13)$$

Now fix $l = 1, \dots, m_i - 1$. For P_π^l -almost all x^l , we have that

$$p_{x^l}(\theta_i) \geq \alpha \quad \text{if and only if} \quad \frac{\pi(\theta_i)}{\alpha} \frac{dP_{\theta_i}^l(x^l)}{dP_\pi^l(x^l)} \geq 1;$$

see Section 1.3. Since $h \geq 0$, we obtain that

$$\begin{aligned} \int_{A^l} h(P_{x^l}) dP_\pi^l(x^l) &\leq \int_{A^l} h(P_{x^l}) \left(\frac{\pi(\theta_i)}{\alpha} \frac{dP_{\theta_i}^l(x^l)}{dP_\pi^l(x^l)} \right) dP_\pi^l(x^l) \\ &= \frac{\pi(\theta_i)}{\alpha} \int_{A^l} h(P_{x^l}) dP_{\theta_i}^l(x^l) \leq \frac{\pi(\theta_i)}{\alpha} \int_{X^l} h(P_{x^l}) dP_{\theta_i}^l(x^l). \end{aligned}$$

Aggregating across l , we deduce that

$$\begin{aligned} h(P_{x^0}) + \sum_{l=1}^{m_i-1} \int_{A^l} h(P_{x^l}) dP_\pi^l(x^l) &\leq h(P_{x^0}) + \frac{\pi(\theta_i)}{\alpha} \sum_{l=1}^{m_i-1} \int_{X^l} h(P_{x^l}) dP_{\theta_i}^l(x^l) \\ &\leq h(\pi, P^{m_i}) + \frac{\pi(\theta_i)}{\alpha} h(\pi, P^{m_i}) \\ &\leq \alpha^2 + \pi(\theta_i)\alpha \leq 2\alpha \end{aligned}$$

where we use the fact that $h(\pi, P^{(n,i)}) \leq \alpha^2$ —see (12). \square

We are ready to conclude the proof. For every $\pi \in \Delta$ and $\alpha \in (0, 1)$, let $\mu_\alpha \in \Delta_\pi^2$ be the random posterior induced by Q^m . Since Δ_π^2 is compact, without loss of generality we can assume that there is $\mu \in \Delta_\pi^2$ such that

$$\lim_{\alpha \rightarrow 0} \mu_\alpha = \mu.$$

For every $\epsilon \in (0, 1)$ and $i = 1, \dots, k - 1$,

$$\mu(\{p : p(\theta_i) \in (\epsilon, 1 - \epsilon)\}) \leq \liminf_{\alpha \rightarrow 0} \mu_\alpha(\{p : p(\theta_i) \in (\epsilon, 1 - \epsilon)\}) = 0$$

where the last equality follows from Claim 2. By continuity, we deduce that

$$\mu(\{p : p(\theta_i) \in (0, 1)\}) = \lim_{\epsilon \rightarrow 0} \mu(\{p : p(\theta_i) \in (\epsilon, 1 - \epsilon)\}) = 0.$$

This shows that

$$\mu = \sum_{i=1}^k \pi(\theta_i) \delta_{\theta_i}^2.$$

Finally, pick $M = 1$ for $n = 1$ and $M = n \cdot (k - 1)$ for $n \geq 2$. It follows from Claim 3 that

$$\lim_{\alpha \rightarrow 0} c_{h^M}(\mu_\alpha) = 0.$$

■

Lemma 7. *If $c_\phi : \Delta^2 \rightarrow [0, \infty)$ is bounded, uniformly posterior-separable, and experimental, then there is $h : \mathcal{E} \rightarrow [0, \infty]$ such that $c_\phi = c_{h^\infty}$.*

Proof. Since c_ϕ is experimental, there must be $n \in \mathbb{N} \cup \{\infty\}$ and $h : \mathcal{E} \rightarrow [0, \infty]$ such that $c_\phi = c_{h^n}$. If $n = \infty$, the desired result follows. Assume therefore that $n < \infty$.

Set $m = 2n$. We claim that for all $\mu \in \Delta^2$,

$$c_{h^n}(\mu) = c_{h^m}(\mu).$$

By definition, $c_{h^n}(\mu) \geq c_{h^m}(\mu)$. Now, fix $\epsilon > 0$. We can find a sequential experiment P^m of length m such that

$$B(\pi, P^m) = \mu \quad \text{and} \quad h(\pi, P^m) \leq c_{h^m}(\mu) + \epsilon.$$

The sequential experiment P^m can be expressed as a series of two sequential experiments $P^{(n,1)}$ and $P^{(n,2)}$, each of length n , such that $P^{(n,2)}$ depends on the outcome of $P^{(n,1)}$. Simple algebra shows that

$$h(\pi, P^m) = h(\pi, P^{(n,1)}) + \int_{X^n} h(p_{x^n}, P_{x^n}^{(n,2)}) dP_\pi^{(n,1)}(x^n)$$

where $p_{x^n} \in \Delta$ is the posterior belief induced by history x^n , prior π , and sequential experiment $P^{(n,1)}$. Since $c_\phi = c_{h^n}$, we obtain that

$$h(\pi, P^m) \geq c_\phi \left(B(\pi, P^{(n,1)}) \right) + \int_{X^n} c_\phi \left(B(p_{x^n}, P_{x^n}^{(n,2)}) \right) dP_\pi^{(n,1)}(x^n).$$

As shown by Bloedel and Zhong (2021),

$$c_\phi \left(B \left(\pi, P^{(n,1)} \right) \right) + \int_{X^n} c_\phi \left(B \left(p_{x^n}, P_{x^n}^{(n,2)} \right) \right) dP_\pi^{(n,1)}(x^n) = c_\phi \left(B \left(\pi, P^m \right) \right).$$

Since $B(\pi, P^m) = \mu$ and $c_\phi = c_{h^n}$, we conclude that

$$c_{h^n}(\mu) \leq c_{h^m}(\mu) + \epsilon.$$

Since the choice of ϵ was arbitrary, we deduce that $c_{h^n}(\mu) \leq c_{h^m}(\mu)$. As mentioned before, $c_{h^n}(\mu) \geq c_{h^m}(\mu)$. We obtain that $c_{h^n}(\mu) = c_{h^m}(\mu)$.

Proceeding by induction, we obtain that for all $k \in \mathbb{N}$, $c_\phi = c_{h^n} = c_{h^{k \cdot n}}$. We conclude that $c_\phi = c_{h^n} = c_{h^\infty}$. \square

Proof of Proposition 1. The “if” direction is trivial. Suppose that c_ϕ is bounded, uniformly posterior-separable, and experimental. By Lemma 7, $c_\phi = c_{h^\infty}$ for some $h : \mathcal{E} \rightarrow [0, \infty]$. Since c_ϕ satisfies FFI, it follows from Proposition 2 that for every $\pi \in \Delta$, there exists a sequence of random posteriors $(\mu_k)_{k=1}^\infty$ in Δ_π^2 such that

$$\lim_{k \rightarrow \infty} \mu_k = \sum_{\theta} \pi(\theta) \delta_{\theta}^2 \quad \text{and} \quad \lim_{k \rightarrow \infty} c_\phi(\mu_k) = 0.$$

By continuity of c_ϕ ,

$$0 = \lim_{k \rightarrow \infty} c_\phi(\mu_k) = c_\phi \left(\lim_{k \rightarrow \infty} \mu_k \right) = c_\phi \left(\sum_{\theta} \pi(\theta) \delta_{\theta}^2 \right).$$

Since c_ϕ is monotone in the convex order, we have that for all $\mu \in \Delta_\pi^2$,

$$c_\phi(\mu) \leq c_\phi \left(\sum_{\theta} \pi(\theta) \delta_{\theta}^2 \right) = 0.$$

We conclude that $c_\phi = 0$. \blacksquare

Proof of Proposition 3. Throughout the proof, we fix a full support prior π^* . For short, we write c_h instead of c_{h^1} .

Claim 4. For all $P \in \mathcal{E}$,

$$h(P) = c_h(B(\pi^*, P)).$$

Proof of the claim. By definition,

$$c_h(B(\pi^*, P)) = \inf\{h(Q) : B(\pi^*, Q) = B(\pi^*, P)\}.$$

If $B(\pi^*, P) = B(\pi^*, Q)$, then $P \sim_b Q$ (see Lemma 2). Thus, since h is Blackwell monotone,

$$\inf\{h(Q) : B(\pi^*, Q) = B(\pi^*, P)\} = h(P).$$

We conclude that $c_h(B(\pi^*, P)) = h(P)$. □

Claim 5. The function c_h is lower semicontinuous on $\Delta_{\pi^*}^2$.

Proof of the claim. Let (μ_n) be a sequence of random posteriors with barycenter π^* that converges to a random posterior μ , which necessarily has barycenter π^* as well. By Lemma 6, there exists a sequence of experiments (P_n) that converges to an experiment P such that (i) for every n , P_n generates μ_n , and (ii) P generates μ . By Claim 4,

$$\liminf_{n \rightarrow \infty} c_h(\mu_n) = \liminf_{n \rightarrow \infty} h(P_n) \quad \text{and} \quad c_h(\mu) = h(P).$$

Since h is lower semicontinuous,

$$\liminf_{n \rightarrow \infty} h(P_n) \geq h(P).$$

We conclude that

$$\liminf_{n \rightarrow \infty} c_h(\mu_n) \geq c_h(\mu).$$

This shows that c_h is lower semicontinuous on $\Delta_{\pi^*}^2$. □

Claim 6. For every $\mu \in \Delta^2$,

$$c_h(\mu) = \inf\{c_h(B(\pi^*, P)) : B(\bar{\mu}, P) \succeq_{cv} \mu\}.$$

Proof of the claim. By definition,

$$c_h(\mu) = \inf\{h(Q) : B(\bar{\mu}, Q) = \mu\} \geq \inf\{h(P) : B(\bar{\mu}, P) \succeq_{cv} \mu\}.$$

Let Q be an experiment that generates μ . Let P be an experiment with $B(\bar{\mu}, P) \succeq_{cv} B(\bar{\mu}, Q)$. By Lemma 2, there exists a stochastic kernel K such that for all $\theta \in \text{supp } \bar{\mu}$, $Q_\theta = KP_\theta$. Define an experiment Q' as follows: for all $\theta \in \Theta$, $Q'_\theta = KP_\theta$. Then, Q' generates μ and $P \succeq_b Q$. Since h is Blackwell monotone, $h(Q) \leq h(P)$. This shows that

$$\inf\{h(Q) : B(\bar{\mu}, Q) = \mu\} \leq \inf\{h(P) : B(\bar{\mu}, P) \succeq_{cv} \mu\}.$$

In sum, we conclude that

$$c_h(\mu) = \inf \{h(P) : B(\bar{\mu}, P) \succeq_{cv} \mu\}.$$

It follows from Claim 4 that

$$c_h(\mu) = \inf \{c_h(B(\pi^*, P)) : B(\bar{\mu}, P) \succeq_{cv} \mu\}.$$

□

Claim 7. The correspondence $\Gamma : \Delta^2 \rightrightarrows \Delta_{\pi^*}^2$ defined by

$$\Gamma(\mu) = \{(B(\pi^*, P)) : B(\bar{\mu}, P) \succeq_{cv} \mu\}$$

has compact values and is upper hemicontinuous.

Proof of the claim. Let (μ_n) be a sequence of random posteriors that converges to a random posterior μ . Let (P_n) be a sequence of experiments such that for every n , $B(\bar{\mu}_n, P_n) \succeq_{cv} \mu_n$. Define $\nu_n = B(\bar{\mu}_n, P_n)$ and $\mu_n^* = B(\pi^*, P_n)$; note that $\mu_n^* \in \Gamma(\mu_n)$. Since $\Delta_{\pi^*}^2$ is compact, without loss of generality we can assume that the sequence of random posteriors (μ_n^*) converges to a random posterior μ^* . Next we show that $\mu^* \in \Gamma(\mu)$.

Let P be an experiment that generates μ^* . Since $\mu_n \rightarrow \mu$, we have that $\bar{\mu}_n \rightarrow \bar{\mu}$. In addition, $B(\pi^*, P_n) \rightarrow B(\pi^*, P)$. It follows from Lemma 4-(i) that $B(\bar{\mu}_n, P_n) \rightarrow B(\bar{\mu}, P)$. Define $\nu = B(\bar{\mu}, P)$, so that $\nu_n \rightarrow \nu$. For every continuous convex function $\phi : \Delta(\Theta) \rightarrow \mathbb{R}$, $\nu_n \succeq_{cv} \mu_n$ implies that

$$\int \phi d\nu_n \geq \int \phi d\mu_n.$$

Since $\nu_n \rightarrow \nu$ and $\mu_n \rightarrow \mu$, we deduce that

$$\int \phi d\nu \geq \int \phi d\mu.$$

This shows that $\nu \succeq_{cv} \mu$, which in turn implies that $\mu^* \in \Gamma(\mu)$. It follows from Aliprantis and Border (2006, Theorem 17.20) that the correspondence Γ has compact values and is upper hemicontinuous. □

We are ready to conclude the proof. By Claim 6,

$$c_h(\mu) = \inf_{\mu^* \in \Gamma(\mu)} c_h(\mu^*).$$

Alternatively,

$$-c_h(\mu) = - \inf_{\mu^* \in \Gamma(\mu)} c_h(\mu^*) = \sup_{\mu^* \in \Gamma(\mu)} -c_h(\mu^*).$$

By Claim 5, the function $-c_h$ is upper semicontinuous on $\Delta_{\pi^*}^2$. By Claim 7, the correspondence Γ is upper hemicontinuous and has (nonempty) compact values. By Aliprantis and Border (2006, Theorem 17.30), the function $-c_h$ is upper semicontinuous on Δ^2 , which means that c_h is lower semicontinuous on Δ^2 . ■

Proof of Proposition 4. Without loss of generality, assume that $\alpha \in (0, 1)$. Define $\pi_\alpha = \alpha\pi + (1 - \alpha)\rho$. Since $\alpha \in (0, 1)$, π and ρ are absolutely continuous with respect to π_α .

First, suppose that $c = c_{h^n}$ for some $n \in \mathbb{N}$. Let P^n be a sequential experiment of length n . Since π and ρ are absolutely continuous with respect to π_α , it follows from Lemma 2 that $B(\pi_\alpha, P^n) = B(\pi_\alpha, P)$ implies $B(\pi, P^n) = B(\pi, P)$ and $B(\rho, P^n) = B(\rho, P)$. In addition,

$$h(\pi_\alpha, P^n) = \alpha h(\pi, P^n) + (1 - \alpha)h(\rho, P^n).$$

We obtain that

$$h(\pi_\alpha, P^n) \geq \alpha c_{h^n}(B(\pi, P)) + (1 - \alpha)c_{h^n}(B(\rho, P)).$$

Since the choice of P^n is arbitrary, we conclude that

$$c_{h^n}(B(\pi_\alpha, P)) \geq \alpha c_{h^n}(B(\pi, P)) + (1 - \alpha)c_{h^n}(B(\rho, P)).$$

Now, suppose that $c = c_{h^\infty}$. Then, we have that

$$\begin{aligned} c_{h^\infty}(B(\pi_\alpha, P)) &= \inf_{n \in \mathbb{N}} c_{h^n}(B(\pi_\alpha, P)) \\ &\geq \inf_{n \in \mathbb{N}} \alpha c_{h^n}(B(\pi, P)) + (1 - \alpha)c_{h^n}(B(\rho, P)) \\ &\geq \alpha \inf_{n \in \mathbb{N}} c_{h^n}(B(\pi, P)) + (1 - \alpha) \inf_{n \in \mathbb{N}} c_{h^n}(B(\rho, P)) \\ &= \alpha c_{h^\infty}(B(\pi, P)) + (1 - \alpha)c_{h^\infty}(B(\rho, P)). \end{aligned}$$

■

Proof of Proposition 5. Define $\phi : \Delta \rightarrow \mathbb{R}$ by

$$\phi(p) = \max_{a \in A} \sum_{\theta} u(a, \theta)p(\theta).$$

Note that ϕ is continuous and convex. Moreover,

$$V(\pi) = \max_{\mu \in \Delta_\pi^2} \int \phi d\mu - c(\mu).$$

Since \mathcal{E} is rich enough, every random posterior is generated by some experiment (see Lemma 1). Thus,

$$V(\pi) = \max_{P \in \mathcal{E}} \int \phi dB(\pi, P) - c(B(\pi, P)).$$

Claim 8. The quantity $\int \phi dB(\pi, P)$ is convex in π .

Proof of the claim. Take $\pi, \rho \in \Delta$ and $\alpha \in [0, 1]$. By Lemma 3,

$$\alpha B(\pi, P) + (1 - \alpha)B(\rho, P) \succeq_{cv} B(\alpha\pi + (1 - \alpha)\rho, P).$$

Define $\mu = B(\pi, P)$ and $\nu = B(\rho, P)$. Since ϕ is convex,

$$\int \phi d(\alpha\mu + (1 - \alpha)\nu) \geq \int \phi dB(\alpha\pi + (1 - \alpha)\rho, P).$$

Since $\int \phi d(\alpha\mu + (1 - \alpha)\nu) = \alpha \int \phi d\mu + (1 - \alpha) \int \phi d\nu$, we obtain that

$$\alpha \int \phi d\mu + (1 - \alpha) \int \phi d\nu \geq \int \phi dB(\alpha\pi + (1 - \alpha)\rho, P).$$

This shows that $\int_\Delta \phi dB(\pi, P)$ is convex in π . □

By hypothesis, $c(B(\pi, P))$ is concave in π . Thus,

$$\int \phi dB(\pi, P) - c(B(\pi, P))$$

is convex in π . Since the supremum of convex functions is a convex function, V is convex. ■

Proof of Proposition 6. (i). Fix a full-support prior π^* . For every experiment $P \in \mathcal{E}$,

$$h_*(P) = c(B(\pi^*, P)) = c_{h^1}(B(\pi^*, P)) \leq h(P).$$

Thus, for every random posterior $\mu \in \Delta^2$,

$$c_{h^1_*}(\mu) \leq c_{h^1}(\mu) = c(\mu). \tag{14}$$

Conversely, for all $P, Q \in \mathcal{E}$ and $\pi \in \Delta$, if $B(\pi^*, P) = B(\pi^*, Q)$, then $B(\pi, P) = B(\pi, Q)$

(see Lemma 2, using the fact that π^* has full support). We obtain that

$$\begin{aligned} c(B(\pi^*, P)) &= c_{h^1}(B(\pi^*, P)) = \inf\{h(Q) : B(\pi^*, Q) = B(\pi^*, P)\} \\ &\geq \inf\{h(Q) : B(\pi, Q) = B(\pi, P)\} \\ &= c_{h^1}(B(\pi, P)) = c(B(\pi, P)). \end{aligned}$$

We deduce that

$$\begin{aligned} c_{h_*^1}(\mu) &= \inf\{h_*(P) : B(\bar{\mu}, P) = \mu\} \\ &= \inf\{c(B(\pi^*, P)) : B(\bar{\mu}, P) = \mu\} \\ &\geq \inf\{c(B(\bar{\mu}, P)) : B(\bar{\mu}, P) = \mu\} = c(\mu). \end{aligned} \tag{15}$$

Combining (14) and (15), we obtain $c = c_{h_*^1}$.

(ii). For every experiment $P \in \mathcal{E}$ and prior $\pi \in \Delta$,

$$c_{h^\infty}(B(\pi, P)) \leq c_{h^1}(B(\pi, P)) \leq h(P).$$

Thus, we have that

$$h_*(P) = \sup_{\pi \in \Delta} c(B(\pi, P)) = \sup_{\pi \in \Delta} c_{h^\infty}(B(\pi, P)) \leq h(P).$$

This shows that for every random posterior $\mu \in \Delta^2$,

$$c_{h_*^\infty}(\mu) \leq c_{h^\infty}(\mu) = c(\mu). \tag{16}$$

Conversely, for every sequential experiment P^n ,

$$\begin{aligned} h_*(\pi, P^n) &= \sum_{i=1}^n \int_{X^{i-1}} h_*(P_{x^{i-1}}) \, dP_\pi^{i-1}(x^{i-1}) \\ &\geq \sum_{i=1}^n \int_{X^{i-1}} c(B(p_{x^{i-1}}, P_{x^{i-1}})) \, dP_\pi^{i-1}(x^{i-1}) \end{aligned}$$

where $p_{x^{i-1}} \in \Delta$ is the posterior belief at partial history x^{i-1} ; by convention, $p_{x^0} = \pi$. As shown by Bloedel and Zhong (2021), since $c = c_{h^\infty}$, we have that

$$\sum_{i=1}^n \int_{X^{i-1}} c(B(p_{x^{i-1}}, P_{x^{i-1}})) \, dP_\pi^{i-1}(x^{i-1}) \geq c(B(\pi, P^n)).$$

Thus, for every random posterior $\mu \in \Delta^2$,

$$c_{h_\star^n}(\mu) = \sup\{h_\star(\bar{\mu}, P^n) : B(\bar{\mu}, P^n) = \mu\} \geq c(\mu).$$

Since the inequality holds for every $n \in \mathbb{N}$, we deduce that

$$c_{h_\star^\infty}(\mu) \geq c(\mu). \quad (17)$$

Combining (16) and (17), we obtain $c_{h_\star^\infty} = c$. ■

Proof of Proposition 7. Matějka and McKay (2015) consider the following information acquisition problem:

$$\max_{P, \beta} \sum_{\theta} \pi(\theta) \int_X \left(\sum_a u(a, \theta) \beta(a|x) \right) dP_{\theta}(x) - c_E(B(\pi, P)). \quad (18)$$

They show that if (P, β) is an optimal solution of (18), then the following conditions hold:

- For all $\theta \in \Theta$ and $a \in A$,

$$\beta_{\theta}(a) = \frac{e^{u(a, \theta)} \beta_{\pi}(a)}{\sum_{a'} e^{u(a', \theta)} \beta_{\pi}(a')}$$

- The probability β_{π} is an optimal solution of

$$\max_{\alpha \in \Delta(A)} \sum_{\theta} \pi(\theta) \log \left(\sum_a e^{u(a, \theta)} \alpha(a) \right). \quad (19)$$

Conversely, if α is an optimal solution of (19), then there is an optimal solution (P, β) of (18) such that $\beta_{\pi} = \alpha$. See Denti, Marinacci and Montrucchio (2020) for an extension of this result to non-discrete environments.

The desired result immediately follows from Matějka and McKay (2015). It is enough to observe that (8) can be rewritten as

$$\max_{P, \beta} \sum_{\theta} \pi^*(\theta) \int_X \left(\sum_a \frac{u(a, \theta) \pi(\theta)}{\pi^*(\theta)} \beta(a|x) \right) dP_{\theta}(x) - c_E(B(\pi^*, P)).$$

Thus, solving (9) is equivalent to solving (8) for a fictitious prior π^* and a fictitious utility $u^* : A \times \Theta \rightarrow \mathbb{R}$ given by $u^*(a, \theta) = u(a, \theta) \pi(\theta) / \pi^*(\theta)$. ■

Proof of Proposition 8. (i). By contradiction, suppose that trade happens with probability one. By (11), the experiment P must be uninformative; otherwise, the buyer would have a profitable deviation (P', β') where P' is uninformative and for every signal x , $\beta'(x) = 1$.

Since P is uninformative, the buyer purchases the good with probability one no matter what the seller's offer is. Thus, σ must put probability one on $s = \max S$; otherwise, the seller would have a profitable deviation in increasing his offer. Since $\max S > v$, the buyer must never purchase the good: $\beta(x) = 0$ for all x . This contradicts the hypothesis that trade happens with probability one. We conclude that in every equilibrium trade fails with positive probability.

(ii). By contradiction, suppose that trade happens with positive probability and the buyer extracts zero surplus. By (11), P must be uninformative; otherwise, the buyer would have a profitable deviation (P', β') where P' is uninformative and for every signal x , $\beta'(x) = 0$. Since P is uninformative, the probability with which the buyer purchases the good does not depend on the seller's offer. Thus, σ must put probability one on $s = \max S$: otherwise, the seller would have a profitable deviation in increasing his offer. Since $\max S > v$, the buyer must never purchase the good: $\beta(x) = 0$ for all x . This contradicts the hypothesis that trade happens with positive probability. We conclude that, in every equilibrium where trade happens with positive probability, the buyer extracts a positive surplus.

(iii). If trade happens with positive probability, then by (ii) the seller must make an offer below v ; otherwise, the buyer's surplus would not be positive. By contradiction, suppose that the seller's offer is never above v : for all $s > v$, $\sigma(s) = 0$. By (11), the experiment P must be uninformative; otherwise, the buyer would have a profitable deviation (P', β') where P' is uninformative and for every signal x , $\beta'(x) = 1$. Since P is uninformative, the probability with which the buyer purchases the good does not depend on the seller's offer. Thus, σ must put probability one on $s = \max S$: otherwise, the seller would have a profitable deviation in increasing his offer. This contradicts the hypothesis that the seller's offer is never above v . We conclude that, in every equilibrium where trade happens with positive probability, the seller randomizes between offers below and above v . ■

C Morris and Strack (2019)

Focusing on a specific environment, Morris and Strack (2019) follow an approach similar to ours. In their paper, the state space is binary: $\Theta = \{\theta_0, \theta_1\}$. The decision maker observes the evolution of a Brownian motion whose drift depends on the state. The flow cost is a linear function of the passage of time, which is of course independent of the evolving beliefs of the decision maker.

Morris and Strack characterize the induced cost on random posteriors $c_{MS} : \Delta^2 \rightarrow [0, \infty]$.

Modulo a scale factor,

$$c_{MS}(\mu) = \phi_{MS}(\bar{\mu}) - \int_{\Delta} \phi_{MS}(p) d\mu(p)$$

with $\phi_{MS}(p) = p(\theta_0) \ln \frac{p(\theta_1)}{p(\theta_0)} + p(\theta_1) \ln \frac{p(\theta_0)}{p(\theta_1)}$.

In particular, c_{MS} is uniformly posterior-separable but not bounded—if $p(\theta_0) \in \{0, 1\}$, then $\phi_{MS}(p) = \infty$.

As observed by Pomatto, Strack and Tamuz (2020), c_{MS} admits a neat representation in terms of Kullback-Leibler divergences (also called relative entropies). Given $\xi_0, \xi_1 \in \Delta(X)$, the *KL-divergence* $D_{KL}(\xi_0 \parallel \xi_1)$ of ξ_1 from ξ_0 is

$$D_{KL}(\xi_0 \parallel \xi_1) = \begin{cases} \int_X \ln(d\xi_1/d\xi_0) d\xi_1 & \text{if } \xi_1 \ll \xi_0, \\ 0 & \text{otherwise.} \end{cases}$$

Simple algebra shows that for each $\pi \in \Delta_+$ and $P \in \mathcal{E}$,

$$c_{MS}(B(\pi, P)) = \pi(\theta_0) D_{KL}(P_{\theta_0} \parallel P_{\theta_1}) + \pi(\theta_1) D_{KL}(P_{\theta_1} \parallel P_{\theta_0}). \quad (20)$$

The quantity $D_{KL}(P_{\theta_0} \parallel P_{\theta_1})$ is a measure of how P_{θ_1} is different from P_{θ_0} (see, e.g., Cover and Thomas, 2012, Chapter 2). Thus, one can view $D_{KL}(P_{\theta_0} \parallel P_{\theta_1})$ as a measure of the informativeness of P when the true state is θ_0 . Analogously, one can view $D_{KL}(P_{\theta_1} \parallel P_{\theta_0})$ as a measure of the informativeness of P when the true state is θ_1 . The cost $c_{MS}(B(\pi, P))$ aggregates, in expectation, these two measures of informativeness.

The results of Morris and Strack (2019) are consistent with our findings. First, c_{MS} does not satisfy FFI. For all $\pi \in \Delta_+$,

$$c_{MS}(\pi(\theta_0)\delta_{\theta_0}^2 + \pi(\theta_1)\delta_{\theta_1}^2) = \infty.$$

As a consequence,

$$\lim_{\pi(\theta_0) \rightarrow 1} c_{MS}(\pi(\theta_0)\delta_{\theta_0}^2 + \pi(\theta_1)\delta_{\theta_1}^2) = \infty > 0.$$

Thus, FFI does not hold. Moreover, c_{MS} is concave in the prior, as easily follows from (20).

In Morris and Strack (2019), the arrival of information is a continuous time process. Bloedel and Zhong (2021) re-derive c_{MS} in a discrete time model with flexible information acquisition. In the language of our paper, Bloedel and Zhong show that $c_{MS} = c_{h\infty}$ for $h : \mathcal{E} \rightarrow [0, \infty]$ given by

$$h(P) = \max \{D_{KL}(P_{\theta_0} \parallel P_{\theta_1}), D_{KL}(P_{\theta_1} \parallel P_{\theta_0})\}.$$

Interestingly, the cost on experiments they use is the same cost on experiments we propose in Proposition 6(ii): it follows from (20) that

$$\max_{\pi \in \Delta} c_{MS}(B(\pi, P)) = \max \{D_{KL}(P_{\theta_0} \| P_{\theta_1}), D_{KL}(P_{\theta_1} \| P_{\theta_0})\}.$$

The cost c_{MS} is consistent with an underlying model of sequential experimentation, but not with an underlying model of one-shot experimentation:

Proposition 9. *There is no cost on experiments $h : \mathcal{E} \rightarrow [0, \infty]$ such that $c_{h^1} = c_{MS}$.*

The result can be readily checked using Proposition 6:

Proof. By contradiction, suppose there is $h : \mathcal{E} \rightarrow [0, \infty]$ such that $c_{h^1} = c_{MS}$. By Proposition 6(i), for all $\pi \in \Delta$ and $P \in \mathcal{E}$,

$$c_{MS}(B(\pi, P)) = \inf \{c_{MS}(B(\pi^*, Q)) : B(\pi, P) = B(\pi, Q)\} \quad (21)$$

where π^* is a fixed full-support prior. The cost c_{MS} is monotone in the convex order. It follows from Lemma 2 that $h_*(P) := c_{MS}(B(\pi^*, P))$ is Blackwell monotone in P . Thus, if π has full support, $B(\pi, P) = B(\pi, Q)$ implies $B(\pi^*, P) = B(\pi^*, Q)$ (by Lemma 2), which in turn implies $h_*(P) = h_*(Q)$, since h_* is monotone in the Blackwell order. We deduce from (21) that for all $\pi \in \Delta_+$ and $P \in \mathcal{E}$,

$$c_{MS}(B(\pi, P)) = c_{MS}(B(\pi^*, P)).$$

Using the representation of c_{MS} in terms of KL-divergence, we obtain that for all $\pi \in \Delta_+$ and $P \in \mathcal{E}$,

$$\pi(\theta_0)D_{KL}(P_{\theta_0} \| P_{\theta_1}) + \pi(\theta_1)D_{KL}(P_{\theta_1} \| P_{\theta_0}) = c_{MS}(B(\pi^*, P)). \quad (22)$$

The KL-divergence is not symmetric. This means that we can find an experiment P such that $D_{KL}(P_{\theta_0} \| P_{\theta_1}) \neq D_{KL}(P_{\theta_1} \| P_{\theta_0})$, that is, we can find an experiment P such that (22) if and only if $\pi = \pi^*$: contradiction. \square

The cost proposed by Morris and Strack (2019) suggests a class of costs on random posteriors that do not satisfy FFI *and* exhibit a non-trivial dependence on prior beliefs. The class is based on a generalization of KL-divergence called f -divergence. Let $f : (0, \infty) \rightarrow \mathbb{R}$ be a convex function such that $f(1) = 0$. Given $\xi_0, \xi_1 \in \Delta(X)$, the f -divergence $D_f(\xi_0 \| \xi_1)$ of ξ_1 from ξ_0 is

$$D_f(\xi_0 \| \xi_1) = \int_X \frac{d\xi_0}{d\lambda} f\left(\frac{d\xi_1/d\lambda}{d\xi_0/d\lambda}\right) d\lambda$$

where λ is a control (σ -finite) measure such that $\xi_0 \ll \lambda$ and $\xi_1 \ll \lambda$.¹² The f -divergences are common generalizations of KL-divergence, used in statistics and information theory (see, e.g., Liese and Vajda, 2006). KL-divergence corresponds to the case in which for all $t \in (0, \infty)$, $f(t) = t \ln t$.

For every $\mu \in \Delta^2$ such that $\bar{\mu} \in \Delta_+$, we define

$$c_f(\mu) = \pi(\theta_0)D_f\left(P_{\theta_0}^\mu\|P_{\theta_1}^\mu\right) + \pi(\theta_1)D_f\left(P_{\theta_1}^\mu\|P_{\theta_0}^\mu\right)$$

where P^μ is an experiment that generates μ ; if $\bar{\mu}(\theta_0) \in \{0, 1\}$, then $c_f(\mu) = 0$. As discussed above for KL-divergence, one can view $D_f(P_{\theta_0}\|P_{\theta_1})$ and $D_f(P_{\theta_1}\|P_{\theta_0})$ as measures of the informativeness of P when the true state is θ_0 and θ_1 , respectively. The cost c_f aggregates, in expectation, these two measures of informativeness.

The function c_f satisfies a number of desirable properties. By standard arguments, c_f is monotone in the convex order and lower semicontinuous. In addition, FFI holds if and only if f is affine (in which case c_f is identical to zero): for all $\pi \in \Delta_+$,

$$c_f(\pi(\theta_0)\delta_{\delta_0} + \pi(\theta_1)\delta_{\delta_1}) = \lim_{t \rightarrow 0} f(t) + \lim_{t \rightarrow 0} \frac{f(t)}{t}.$$

The right-hand side is equal to zero if and only if f is affine. Thus, if f is not affine,

$$\lim_{\pi(\theta_0) \rightarrow 1} c_f(\pi(\theta_0)\delta_{\delta_0} + \pi(\theta_1)\delta_{\delta_1}) = \lim_{t \rightarrow 0} f(t) + \lim_{t \rightarrow 0} \frac{f(t)}{t} > 0,$$

which shows that FFI does not hold. Reasoning as above, one can also check that $c_f(B(\pi, P))$ is concave in π ; one can also show that $c_f(B(\pi, P))$ is convex in P . An interesting question for future research is whether c_f is experimental.

D Bloedel and Zhong (2021), Hébert and Woodford (2021)

A few recent papers have explored the relationship between rational inattention and sequential information acquisition. Bloedel and Zhong (2021) and Hébert and Woodford (2021) address questions similar to ours, but allow the flow cost of information to depend arbitrarily on the evolving beliefs of the decision maker, as in rational inattention. As a consequence, more cost functions on random posteriors can be generated; the next example is a concrete illustration.

Example 1 (Bloedel and Zhong, 2021). Let $\phi : \Delta \rightarrow \mathbb{R}$ be a concave continuous function. In contrast with (2), suppose the decision maker incurs a cost $c_\phi(B(p_{x^{i-1}}, P_{x^{i-1}}))$ for per-

¹²Behind the integral the conventions $f(0) = \lim_{t \rightarrow 0} f(t)$ and $0 \cdot f(t/0) = t \lim_{s \rightarrow \infty} f(s)/s$ are adopted.

forming an experiment $P_{x^{i-1}} \in \mathcal{E}$ at a history $x^{i-1} \in X^{i-1}$, where $p_{x^{i-1}} \in \Delta$ is the posterior belief conditional on x^{i-1} (by convention, $p_{x^0} = \pi$).

Given a prior $\pi \in \Delta$, the resulting cost of a sequential experiment P^n is

$$h(\pi, P^n) = \sum_{i=1}^n \int_{X^{i-1}} c_\phi(B(p_{x^{i-1}}, P_{x^{i-1}})) \, dP_\pi^{i-1}(x^{i-1}).$$

As Bloedel and Zhong (2021) show, $c_\phi(B(\pi, P^n)) = h(\pi, P^n)$. We obtain that for all $\mu \in \Delta^2$,

$$c_\phi(\mu) = \inf\{h(\pi, P^n) : B(\bar{\mu}, P^n) = \mu\}.$$

▲

Thus, if the flow cost depends arbitrarily on the evolving beliefs of the decision maker, then *any* uniformly posterior-separable cost function can be generated—e.g., the entropy-based cost of Matějka and McKay (2015). Our Proposition 1 adds a caveat: the arbitrariness is crucial; if the flow cost depends only on the per-period experiment, then *no* cost function that is bounded and uniformly posterior-separable can be generated.

As a special case of their framework, Bloedel and Zhong (2021) study what happens when the flow cost depends only on the per-period experiment, as in (2). Their findings are consistent with ours. In particular, under the hypothesis that h is “locally quadratic”—we refer the reader to their paper for a precise definition of this property—they show that no bounded, non-trivial, uniformly posterior-separable cost is consistent with a primitive model of sequential information acquisition (Bloedel and Zhong, 2021, Proposition 3).¹³ Our Proposition 2 complements their result, as we put no functional-form assumption on h , generalize beyond uniform posterior separability, and rely on a different argument for the proof. As discussed above, Proposition 2 generalizes to the case in which the flow cost of information depends on the history of experimentation, a case that is not discussed by Bloedel and Zhong (2021).

E Expertise and fatigue

In the main text, we consider models of sequential information acquisition where the flow cost of information does not depend on the history of experimentation. This invariance property rules out interesting phenomena such as building expertise and experiencing fatigue (see, e.g., Dillenberger, Krishna and Sadowski, 2022). Next we discuss how to incorporate these phenomena into our analysis and results.

¹³When there are at least three states, their result holds for all non-trivial uniformly posterior-separable costs, bounded or unbounded.

We assume that the decision maker incurs a cost $h(P_{x^{i-1}}|P_{x^0}, \dots, P_{x^{i-2}})$ for performing an experiment $P_{x^{i-1}}$ at a history $x^{i-1} = (x_0, \dots, x_{i-2})$. The cost may depend on the experiments $P_{x^0}, \dots, P_{x^{i-2}}$ that the decision maker performed in the past. Building expertise corresponds to the case in which

$$h(P_{x^{i-1}}|P_{x^0}, \dots, P_{x^{i-2}}) \leq h(P_{x^{i-1}}), \quad (23)$$

that is, the more information the agent acquires, the less costly information becomes. Experiencing fatigue is the opposite phenomenon.

Given prior $\pi \in \Delta$, the expected cost of a sequential experiment P^n is

$$h(\pi, P^n) = \sum_{i=1}^n \int_{X^{i-1}} h(P_{x^{i-1}}|P_{x^0}, \dots, P_{x^{i-2}}) dP_\pi^{i-1}(x^{i-1})$$

where $P_\pi^{i-1} \in \Delta(X^{i-1})$ is the predictive probability generated by π and the first $i-1$ experiments.

When the state is binary, Proposition 2 and its proof extend verbatim to the case in which the flow cost of information depends on the history of experimentation, regardless of the nature of the dependence. When there are many states, Proposition 2 extends verbatim to the case in which the decision maker builds expertise, as defined by (23); the proof needs only a minor adjustment: replace (13) with

$$\begin{aligned} \sum_{l=1}^{m_i-1} \int_{B^l} h(p_{x^l}, Q^{(m,i+1)}|P_{x^0}, \dots, P_{x^{l-1}}) dP_\pi^l(x^l) &\leq \sum_{l=1}^{m_i-1} \int_{B^l} h(p_{x^l}, Q^{(m,i+1)}) dP_\pi^l(x^l) \\ &\leq 2(k-i-1)\alpha \end{aligned}$$

where the first inequality uses the fact that the agent builds expertise. Intuitively, building expertise is the most interesting case for FFI, since the property seems to suggest that the more the decision maker learns, the less costly information becomes.

Proposition 4 also holds whether or not the flow cost depends on the history of experimentation. Proposition 6-(ii) instead is specific to the case in which

$$h(P_{x^{i-1}}|P_{x^0}, \dots, P_{x^{i-2}}) = h(P_{x^{i-1}}).$$

Proposition 8 generalizes under the mild assumption that running an uninformative experiment today does not reduce to zero the cost of running an informative experiment tomorrow.

F Additional examples on continuity

Example 2. Let $\Theta = \{Low, High\}$ and $X = [0, 1]$. Let P_0 be an uninformative experiment that always generates the signal $x = 0$. For every $k \in \mathbb{N}$, let P_k be an experiment that generates the signal $x = 0$ with probability one when $\theta = Low$, and that generates the signal $x = 1/k$ with probability $1 - 1/k$ and the signal $x = 0$ with the remaining probability $1/k$ when $\theta = High$. Notice that as $k \rightarrow \infty$, $P_k \rightarrow P_0$.

Assume that $h(P_0) = 0$ and for every $k \in \mathbb{N}$, $h(P_k) = 0$. For every other experiment P , let $h(P) = 1$. The cost $h : \mathcal{E} \rightarrow \{0, 1\}$ is lower semicontinuous, but not Blackwell monotone. For example, let Q_k be an experiment that generates the signal $x = 0$ with probability one when $\theta = Low$, and that generates the signal $x = 1$ with probability $1 - 1/k$ and the signal $x = 0$ with the remaining probability $1/k$ when $\theta = High$. Clearly, $Q_k \sim_b P_k$. However, $h(Q_k) = 1 > h(P_k) = 0$.

The experimental cost c_{h^1} is not lower semicontinuous. For example, assuming a uniform prior, let μ_k be the random posterior generated by the experiment P_k , and let μ be the random posterior corresponding to full information. It is easy to check that $\mu_k \rightarrow \mu$, $c_{h^1}(\mu_k) = 0$ for all k , but $c_{h^1}(\mu) = 1$. This shows that c_{h^1} is not lower semicontinuous. \blacktriangle

Example 3. Let π^* be the uniform prior. Define a cost on experiments $h : \mathcal{E} \rightarrow [0, \infty)$ by

$$h(P) = \sum_{\theta \in \Theta} \pi^*(\theta) D_{KL}(P_\theta \| P_{\pi^*})$$

where D_{KL} is the Kullback-Leibler divergence. It is easy to check—recall that another name for the KL divergence is relative entropy—that

$$h(P) = c_E(B(\pi^*, P)).$$

Thus, $h(P)$ is the expected reduction in entropy with respect to a reference prior π^* —see also (7) in the main text. By standard arguments, h is Blackwell monotone and lower semicontinuous.

Define $c : \Delta^2 \rightarrow [0, \infty)$ by

$$c(\mu) = \inf \{h(P) : B(\bar{\mu}, P) = \mu\}.$$

By definition, c is experimental and non-trivial. By Proposition 3, c is lower semicontinuous.

Next we verify that c is continuous on each subdomain Δ_π^2 . But first we need a closed-form characterization of $c(\mu)$:

Claim 9. For every $\mu \in \Delta^2$,

$$c(\mu) = \pi^*(T) \sum_{\theta \in T} \pi_T^*(\theta) D_{KL} \left(P_\theta^\mu \| P_{\pi_T^*}^\mu \right)$$

where T is the support of $\bar{\mu}$, π_T^* is the uniform distribution over T , and P^μ is an experiment that generates μ .

Proof of the claim. Define experiments Q and Q' by

$$Q_\theta = \begin{cases} P_\theta^\mu & \text{if } \theta \in T \\ P_{\pi_T^*}^\mu & \text{if } \theta \notin T \end{cases} \quad \text{and} \quad Q'_\theta = \begin{cases} P_\theta^\mu & \text{if } \theta \notin T \\ P_{\pi_{T'}^*}^\mu & \text{if } \theta \in T \end{cases}$$

where $T' = \Theta \setminus T$. An application of the chain rule (see, e.g., Cover and Thomas, 2012, Section 2.5) shows that

$$h(P) = \pi^*(T) D(Q_{\pi^*} \| P_{\pi^*}^\mu) + \pi^*(T') D(Q'_{\pi^*} \| P_{\pi^*}^\mu) + h(Q) + h(Q')$$

where we use the fact that $Q_{\pi^*} = P_{\pi_T^*}^\mu$ and $Q'_{\pi^*} = P_{\pi_{T'}^*}^\mu$. Thus,

$$\begin{aligned} h(P^\mu) &\geq h(Q) = \pi^*(T) \sum_{\theta \in T} \pi_T^*(\theta) D_{KL} (Q_\theta^\mu \| Q_{\pi_T^*}^\mu) \\ &= \pi^*(T) \sum_{\theta \in T} \pi_T^*(\theta) D_{KL} \left(P_\theta^\mu \| P_{\pi_T^*}^\mu \right). \end{aligned}$$

Since for all $\theta \in T$, $Q_\theta = P_\theta^\mu$, the experiment Q generates μ . We obtain that

$$c(\mu) \leq \pi^*(T) \sum_{\theta \in T} \pi_T^*(\theta) D_{KL} \left(P_\theta^\mu \| P_{\pi_T^*}^\mu \right).$$

If Q^μ is any other experiment that generates μ , then there exists a stochastic kernel K such that for all $\theta \in T$, $Q_\theta = K Q_\theta^\mu$ (see Lemma 2). Thus,

$$\begin{aligned} \sum_{\theta \in T} \pi_T^*(\theta) D_{KL} \left(Q_\theta^\mu \| Q_{\pi_T^*}^\mu \right) &\geq \sum_{\theta \in T} \pi_T^*(\theta) D_{KL} \left(Q_\theta \| Q_{\pi_T^*} \right) \\ &= \sum_{\theta \in T} \pi_T^*(\theta) D_{KL} (Q_\theta \| Q_{\pi^*}). \end{aligned}$$

We conclude that

$$c(\mu) = \pi^*(T) \sum_{\theta \in T} \pi_T^*(\theta) D_{KL} \left(P_\theta^\mu \| P_{\pi_T^*}^\mu \right).$$

□

Now we are ready to prove that c is continuous on each subdomain Δ_π^2 . Fix a prior $\pi \in \Delta$ and let T be its support. Simple algebra shows that

$$\sum_{\theta \in T} \pi_T^*(\theta) D_{KL} \left(P_\theta^\mu \| P_{\pi_T^*}^\mu \right) = c_E(B(\pi_T^*, P^\mu)).$$

It follows from Claim 9 that

$$c(\mu) = \pi^*(T) c_E(B(\pi_T^*, P^\mu)).$$

Now, let (μ_n) be a sequence of random posteriors—all with barycenter π —that converges to μ . By Lemma 4, $B(\pi_T^*, P^{\mu_n}) \rightarrow B(\pi_T^*, P^\mu)$. By continuity of c_E , we deduce that $c(\mu_n) \rightarrow c(\mu)$. This shows that c is continuous on each subdomain Δ_π^2 . \blacktriangle

G Ultimatum game: a tractable functional form

In the main text, we study the ultimatum game under broad conditions on the cost over experiments. In this section, we provide an explicit characterization of the equilibria for the tractable functional form we introduce in Section 2.

We fix a full-support probability distribution over offers $\pi^* \in \Delta(S)$ —for simplicity, let π^* be the uniform distribution. We consider a cost over experiments h_k by

$$h_k(P) := k \cdot h_{E^*}(P) = k \cdot c_E(B(\pi^*, P)). \quad (24)$$

Thus, h_k is the expected reduction in the entropy of beliefs *with respect to a reference prior* π^* , scaled by a constant $k > 0$ that parameterizes the marginal cost of information. Since π^* has full support, $c_E(B(\pi^*, P)) = 0$ if and only if P is uninformative. As a consequence, (11) holds.

We distinguish between two cases. First, we consider a simple set of offers that includes only three options; the simplicity of the setup allows us to highlight the main features of the game. Next, we consider richer sets of offers that approximate the continuum of offers considered by Ravid (2020). Throughout, let

$$\beta_s := \int_X \beta(x) dP_s(x)$$

be the probability of trade when the offer is s , and let $\beta_\sigma := \sum_s \sigma(s) \beta_s$ be the marginal probability of trade.

G.1 Simple set of offers

We assume that $S = \{2v/3, v, 2v\}$: the seller can offer the good at two-thirds of its value, its value, or twice its value. The next proposition characterizes the equilibria where trade happens with positive probability.

Proposition 10. *Assume (24) for the cost of information, and let $S = \{2v/3, v, 2v\}$. A strategy profile (σ, P, β) is an equilibrium with $\beta_\sigma > 0$ if and only if there exists $z \in (0, 2v/3)$ that satisfies the following conditions:*

(i) For all $s \in S$,

$$\beta_s = \frac{z}{s}$$

(ii) For all $s \neq v$,

$$\sigma(s) = \frac{k/3}{v-s} \log \frac{v-z}{s-z}.$$

In particular, an equilibrium where trade happens exists if and only if

$$k < \frac{3v}{3 \log 3 - 2 \log 2}.$$

In every equilibrium where trade occurs, the seller is indifferent between all offers: for all $s, s' \in S$,

$$s\beta_s = z = s'\beta_{s'}. \quad (25)$$

The seller randomizes over offers to make it optimal for the buyer to choose a strategy (P, β) such that, indeed, (25) holds. In particular, the seller always puts positive probability on $s = 2v/3$ and $s = 2v$; this creates an incentive for the buyer to acquire information. Finally, there exists an equilibrium where trade occurs if and only if the marginal cost of information k is small relative to the value v of the good; otherwise, the buyer's incentive to acquire information is not strong enough to monitor the seller's offer and she chooses never to buy the good. By an argument analogous to Ravid (2020, p. 2959), the buyer earns a positive payoff (consumer surplus minus attention costs).

The proof of Proposition 10 builds on the logit characterization of optimal information acquisition provided by Proposition 7.

Proof. Let (σ, P, β) be an equilibrium with $\beta_\sigma > 0$. By Proposition 8, $\sigma(2v) > 0$ and $\sigma(\frac{2v}{3}) > 0$. The seller's indifference condition between $s = \frac{2v}{3}$ and $s = 2v$ is

$$\frac{2v}{3}\beta_{\frac{2v}{3}} = 2v\beta_{2v} \quad \iff \quad \beta_{\frac{2v}{3}} = 3\beta_{2v}.$$

From Proposition 7, $\beta_v = \beta_{\sigma^*}$. Thus

$$\frac{1}{3}\beta_{\frac{2v}{3}} + \frac{1}{3}\beta_v + \frac{1}{3}\beta_{2v} = \beta_{\sigma^*} \quad \Longleftrightarrow \quad \frac{1}{2}\beta_{\frac{2v}{3}} + \frac{1}{2}\beta_{2v} = \beta_{\sigma^*}$$

It follows from $\beta_{\frac{2v}{3}} = 3\beta_{2v}$ that

$$\beta_{\frac{2v}{3}} = \frac{3}{2}\beta_{\sigma^*} \quad \text{and} \quad \beta_{2v} = \frac{1}{2}\beta_{\sigma^*}.$$

In particular, the seller is indifferent between all offers. From Proposition 7,

$$\frac{e^{\frac{(v-\frac{2v}{3})\sigma(\frac{3v}{3})}{k/3}}\beta_{\sigma^*}}{e^{\frac{(v-\frac{2v}{3})\sigma(\frac{2v}{3})}{k/3}}\beta_{\sigma^*} + 1 - \beta_{\sigma^*}} = \frac{3}{2}\beta_{\sigma^*} \quad \text{and} \quad \frac{e^{\frac{(v-2v)\sigma(2v)}{k/3}}\beta_{\sigma^*}}{e^{\frac{(v-2v)\sigma(2v)}{k/3}}\beta_{\sigma^*} + 1 - \beta_{\sigma^*}} = \frac{1}{2}\beta_{\sigma^*}.$$

Simple algebra shows that

$$e^{\frac{(v-\frac{2v}{3})\sigma(\frac{3v}{3})}{k/3}} \left(1 - \frac{3}{2}\beta_{\sigma^*}\right) = \frac{3}{2}(1 - \beta_{\sigma^*}) \quad \text{and} \quad e^{\frac{(v-2v)\sigma(2v)}{k/3}} \left(1 - \frac{1}{2}\beta_{\sigma^*}\right) = \frac{1}{2}(1 - \beta_{\sigma^*}).$$

Thus $\beta_{\sigma^*} \in (0, \frac{2}{3})$. Simplifying further, we obtain

$$\sigma\left(\frac{2v}{3}\right) = \frac{k/3}{v/3} \log \frac{3(1 - \beta_{\sigma^*})}{2 - 3\beta_{\sigma^*}} \quad \text{and} \quad \sigma(2v) = \frac{k/3}{v} \log \frac{2 - \beta_{\sigma^*}}{1 - \beta_{\sigma^*}}.$$

Taking $z = v\beta_{\sigma^*}$, we observe that (i) and (ii) hold.

Now take $z \in (0, 2v/3)$ such that (i) and (ii) hold. Notice that $\beta_{\sigma^*} = z/v$. Simple algebra shows that

$$\frac{1}{3} \sum_s \frac{e^{\frac{(v-s)\sigma(s)}{k/3}}}{e^{\frac{(v-s)\sigma(s)}{k/3}} \frac{z}{v} + 1 - \frac{z}{v}} = 1.$$

That is the first order condition of the auxiliary maximization problem in Proposition 7.

Thus we can find a best response (P, β) to σ such that $\beta_{\sigma^*} = \frac{z}{v}$. In addition, for all $t \in T$,

$$\beta_s = \frac{e^{\frac{(v-s)\sigma(s)}{k/3}} \frac{z}{v}}{e^{\frac{(v-s)\sigma(s)}{k/3}} \frac{z}{v} + 1 - \frac{z}{v}} = \frac{z}{s}.$$

Therefore the seller is indifferent between all offers: σ is a best response to (P, β) as well.

We conclude that (σ, P, β) is an equilibrium with $\beta_{\sigma} > 0$.

Finally, consider the function $f : (0, \frac{2v}{3}) \rightarrow [0, \infty)$ given by

$$f(z) = \frac{k}{v} \log \frac{3(v-z)}{2v-3z} + \frac{k}{3v} \log \frac{2v-z}{v-z}.$$

Note that f is continuous and strictly increasing. In addition,

$$\lim_{z \rightarrow 0} f(z) = \frac{k}{v} \log \frac{3}{2} + \frac{k}{3v} \log 2 \quad \text{and} \quad \lim_{z \rightarrow \frac{2v}{3}} f(z) = \infty.$$

As shown above, if (σ, P, β) is an equilibrium with $\beta_\sigma > 0$, then

$$\lim_{z \rightarrow 0} f(z) < f(v\beta_{\sigma^*}) = \sigma \left(\frac{2}{3}v \right) + \sigma(2v) \leq 1.$$

Conversely, if $\lim_{z \rightarrow 0} f(z) < 1$, then by Brouwer's fixed point theorem there exists z such that $f(z) \leq 1$. As shown above, this implies that there exists an equilibrium (σ, P, β) with $\beta_\sigma > 0$. The condition $\lim_{z \rightarrow 0} f(z) < 1$ is equivalent to $k < \frac{3v}{3 \log 3 - 2 \log 2}$. \square

G.2 Rich sets of offers

Next, we generalize Proposition 10 to rich sets of offers that approximate the continuum of offers studied by Ravid (2020).

For ease of exposition, assume $v = 1$: it is common knowledge that the good has value 1. We consider a finite set of offers S that satisfies the following property: there exists $\hat{s} \in S$ with $0 < \hat{s} < 1$ such that

$$\sum_{s \geq \hat{s}} \frac{1}{s} = |\{s \in S : s \geq \hat{s}\}|. \quad (26)$$

For example, $S = \{2/3, 1, 2\}$ and $\hat{s} = 2/3$, as in Proposition 10. Next is a more interesting example, which shows how to construct finer and finer sets of offers that satisfy (26).

Example 4. Take any increasing sequence $0 \leq s_1 < \dots < s_n \leq 1$. Choose an index $m \in \{1, \dots, n\}$ such that $s_m \in (1/2, 1)$. Define

$$S = \left\{ s_1, \dots, s_n, \frac{s_n}{2s_n - 1}, \dots, \frac{s_m}{2s_m - 1} \right\}.$$

For $\hat{s} = s_m$, condition (26) holds. ▲

We impose an addition condition on S . To state it, define

$$\hat{S} = \{s \in S : s \geq \hat{s} \text{ and } s \neq 1\}.$$

We assume that there exists $z \in (0, \hat{t})$ such that

$$\sum_{s \in \hat{S}} \frac{k}{(1-s)|S|} \log \frac{1-z}{s-z} = 1 \quad (27)$$

where $k > 0$ is the scale factor that parametrizes the marginal cost of information. The next lemma provides a simplification:

Lemma 8. *A sufficient condition for (27) is that*

$$\frac{|\hat{S}| \log \hat{s}}{|S| \hat{s} - 1} < \frac{1}{k}. \quad (28)$$

In addition,

$$z \geq \frac{1 - \hat{s} \exp\left(\frac{|S| \cdot (1-\hat{s})}{|\hat{S}| \cdot k}\right)}{1 - \exp\left(\frac{|S| \cdot (1-\hat{s})}{|\hat{S}| \cdot k}\right)} > 0.$$

Proof. For $z \in (0, \hat{t})$ and $s \in [\hat{s}, \infty)$, define

$$f(z, s) = \frac{k}{(1-s)|S|} \log \frac{1-z}{s-z}.$$

The quantity $f(z, s)$ is increasing in z and decreasing in s :

$$\frac{\partial}{\partial z} f(z, s) > 0 \quad \Leftrightarrow \quad \frac{1}{(s-z)(1-z)} > 0.$$

$$\frac{\partial}{\partial s} f(z, s) < 0 \quad \Leftrightarrow \quad \log \frac{1-z}{s-z} < \frac{1-s}{s-z} = \frac{1-z}{s-z} - 1.$$

Define $F(z) = \sum_{s \in \hat{S}} f(z, s)$. Notice that

$$\lim_{z \nearrow \hat{s}} F(z) = \lim_{z \nearrow \hat{s}} f(z, \hat{s}) = \infty.$$

Moreover, being $f(z, s)$ decreasing in s ,

$$\lim_{z \searrow 0} F(z) \leq |\hat{S}| \lim_{z \searrow 0} f(z, \hat{s}) = k \frac{|\hat{S}| \log \hat{s}}{|S| \hat{s} - 1} < 1$$

where the last inequality follows from (28). Since $F(z)$ is continuous in z , there must exist z such that $F(z) = 1$.

Define

$$\hat{z} = \frac{1 - \hat{s} \exp\left(\frac{|S| \cdot (1-\hat{s})}{|\hat{S}| \cdot k}\right)}{1 - \exp\left(\frac{|S| \cdot (1-\hat{s})}{|\hat{S}| \cdot k}\right)}.$$

Observe that $|\hat{S}| \cdot f(\hat{z}, \hat{s}) = 1$. Thus, $\hat{z} \in (0, \hat{s})$. Since $f(z, s)$ is decreasing in s ,

$$F(\hat{z}) \leq |\hat{S}| \cdot f(\hat{z}, \hat{s}) = 1.$$

Since $f(z, s)$ is increasing in z , $F(z) = 1$ implies $z \geq \hat{z}$. □

Next we provide an example of a rich set of offers that satisfy both (26) and (27).

Example 5. Fix $n \geq 2$. Define

$$S = \left\{ 0, \frac{1}{2^n}, \dots, \frac{2^n - 1}{2^n}, 1, \frac{2^n - 1}{2(2^n - 1) - 2^n}, \dots, \frac{3}{2} \right\}.$$

For $\hat{s} = 3/4$, (26) holds. For n sufficiently large,

$$\frac{|\hat{S}|}{|S|} \approx \frac{2}{5}.$$

Condition (28) becomes

$$k < \frac{5}{8(\log 4 - \log 3)}.$$

Thus, (27) is satisfied as long as the marginal cost of information is sufficiently small. ▲

We are now ready to generalize Proposition 10.

Proposition 11. *Assume (24) for the cost of information, and let S satisfy (26) and (27). For $z \in (0, \hat{t})$ that satisfies (27), there exists an equilibrium (σ, P, β) with the following properties:*

- The support of σ is \hat{S} .

- For all $s \in \hat{S}$,

$$\sigma(s) = \frac{k/|S|}{1-s} \ln \frac{1-z}{s-z}.$$

- For all $s < \hat{s}$, $\beta_s = z$.

- For all $s \geq \hat{s}$, $\beta_s = z/s$.

What happens as the set of offers become finer and finer? In the equilibrium described above, $\sigma(s)$ is decreasing in $s \in \hat{S}$. It follows that

$$\begin{aligned} \sigma(\{s' : s' \leq s\}) &\geq |\hat{S} \cap [0, s]| \cdot \sigma(s) \\ &= \frac{|\hat{S} \cap [0, s]|}{|S|} \frac{k}{1-s} \log \frac{1-z}{s-z}. \end{aligned}$$

Now imagine that the set of possible offers become finer and finer. If \hat{s} is fixed and the ratio $|\hat{S} \cap [0, s]|/|S|$ is bounded away from zero, then the probability that the seller claims less than s is bounded away from zero, being z bounded away from zero (cfr. Lemma 8). Thus, in particular, σ does *not* converge in distribution to a Dirac measure concentrated on 1. This shows that, even in the limit of a continuum of offers, the seller randomizes between offers above and below the value of the good.

We can repeat the same exercise for the probability of trade. Being $\sigma(s)$ decreasing $s \in \hat{S}$, we have

$$\begin{aligned} \sum_{s \in \hat{S}} \beta_s \sigma(s) &\geq |\hat{S}| \beta_{\max(S)} \sigma(\max(S)) \\ &= \frac{|\hat{S}|}{|S|} \frac{z}{\max(S)} \frac{k}{(1 - \max(S))} \log \frac{1 - z}{\max(S) - z}. \end{aligned}$$

Now imagine that the set of possible offers become finer and finer. If \hat{s} and $\max(S)$ are fixed and the ratio $|\hat{S}|/|S|$ is bounded away from zero, then the probability of trade is bounded away from zero, being z bounded away from zero (cfr. Lemma 8). This shows that, even in the limit of a continuum of offers, trade happens with positive probability.

Example 5 (Continued). Let n be large. The distance between two consecutive points s_i and s_{i+1} in $\hat{S} \cap [0, 1]$ is $1/2^n$. Thus, for $s \in [3/4, 1]$,

$$\begin{aligned} \sigma(\{s' : s' \leq s\}) &= \sum_{s' \in \hat{S} \cap [0, 1]} \frac{k/|S|}{1 - s'} \log \frac{1 - z}{s' - z} \\ &= \frac{2^n}{|S|} \sum_{s' \in \hat{S} \cap [0, 1]} \frac{1}{2^n} \frac{k}{1 - s'} \log \frac{1 - z}{s' - z} \\ &\approx \frac{3}{5} \int_{\frac{3}{4}}^s \frac{k}{1 - s'} \log \frac{1 - z}{s' - z} ds' \end{aligned}$$

where the last approximation presumes that n is large.

The distance between two consecutive points s_i and s_{i+1} in $\hat{S} \cap (1, \infty)$ is

$$\frac{(2s_i - 1)(2s_{i+1} - 1)}{2^n}.$$

Thus, for $s \in (1, 3/2]$,

$$\begin{aligned}\sigma(\{s' : 1 < s' \leq s\}) &= \sum_{s' \in \hat{S} \cap (1, \infty)} \frac{k/|S|}{1-s'} \log \frac{1-z}{s'-z} \\ &\approx \frac{3}{5} \int_{\frac{3}{4}}^s \frac{(2s'-1)^2}{1-s'} \log \frac{1-z}{s'-z} ds'\end{aligned}$$

where the last approximation presumes that n is large.

Overall, in the limit, the equilibrium is as follows:

- σ has support $[3/4, 3/2]$ and admits a density f given by

$$f(s) = \begin{cases} \frac{3k}{5} \frac{1}{1-s} \log \frac{1-z}{s-z} & s \in [3/4, 1] \\ \frac{3k}{5} \frac{(2s-1)^2}{1-s} \log \frac{1-z}{s-z} & s \in (1, 3/2]. \end{cases}$$

- z is pinned down by the equation

$$\int_{\frac{3}{4}}^1 \frac{k}{1-s} \log \frac{1-z}{s-z} ds + \int_1^{\frac{3}{2}} \frac{k}{1-s} \log \frac{1-z}{s-z} ds = \frac{5}{3}.$$

- For all $s < 3/4$, $\beta_s = z$.
- For all $s \geq 3/4$, $\beta_s = z/s$.

▲

We conclude with the proof of Proposition 11, which follows the same steps of the proof of Proposition 10.

Proof of Proposition 11. For all $s \geq \hat{s}$,

$$s\beta_s = z = \beta_1.$$

Thus the seller is indifferent among all $s \geq \hat{s}$. For all $s < \hat{s}$,

$$s\beta_s = sz < z = \beta_1$$

where we use the fact that $\hat{s} < 1$. Thus, σ is a best reply to (P, β) .

It follows from (26) that

$$\beta_{\sigma^*} = \sum_{s \in S} \frac{1}{|S|} \beta_s = z$$

Thus, from Proposition 7 in the main text, (P, β) is a best reply to σ if for all $s \in S$,

$$\beta_s = \frac{e^{\frac{(1-s)\sigma(s)}{k/|S|}} z}{e^{\frac{(1-s)\sigma(s)}{k/|S|}} z + 1 - z}.$$

If $s < \hat{s}$, then $\sigma(s) = 0$, which implies

$$\frac{e^{\frac{(1-s)\sigma(s)}{k/|S|}} z}{e^{\frac{(1-s)\sigma(s)}{k/|S|}} z + 1 - z} = z.$$

If $s = 1$, then (regardless of σ)

$$\frac{e^{\frac{(1-s)\sigma(s)}{k/|S|}} z}{e^{\frac{(1-s)\sigma(s)}{k/|S|}} z + 1 - z} = z.$$

If $s \in \hat{S}$, the choice of $\sigma(s)$ guarantees that that

$$\frac{e^{\frac{(1-s)\sigma(s)}{k/|S|}} z}{e^{\frac{(1-s)\sigma(s)}{k/|S|}} z + 1 - z} = \frac{z}{s}.$$

This shows that (P, β) is a best reply to σ . We conclude that (σ, P, β) is an equilibrium. \square

H Costly monitoring in leader-follower games

The inconsistency of posterior-based costs with a primitive model of costly experimentation has concrete implications for the analysis of information acquisition in games. In the main text, we make the point by focusing on an ultimatum game proposed by Ravid (2020). Here we show that the same conclusions hold for any leader-follower game.

There are two players in the game, a leader and a follower. The leader chooses an action a from a finite set A , the follower chooses an action b from a finite set B ; denote by $u_l(a, b)$ and $u_f(a, b)$ the corresponding utilities. We make the generic assumption that utilities are different across actions: for all $a, a' \in A$ and $b, b' \in B$ such that $a \neq a'$ and $b \neq b'$,

$$u_l(a, b) \neq u_l(a', b) \quad \text{and} \quad u_f(a, b) \neq u_f(a, b'). \quad (29)$$

As usual, for $\alpha \in \Delta(A)$ and $\beta \in \Delta(B)$, we denote by $u_l(\alpha, \beta)$ and $u_f(\alpha, \beta)$ the mixed extensions of the utilities.

Before taking action, the follower has the opportunity to monitor the sender's action at a cost. A monitoring structure is represented by an experiment $P : A \rightarrow \Delta(X)$. A strategy

for the leader is a probability over actions $\alpha \in \Delta(A)$. A strategy for the follower consists of an experiment $P : A \rightarrow \Delta(X)$ and a measurable action rule $\sigma : X \rightarrow \Delta(B)$ that specifies for every signal x , the probability $\sigma(b|x)$ with which the follower takes an action b .

Next we analyze the game for two different specifications of the cost of information, one that takes as primitive a cost on random posteriors, as in applications of rational inattention, and one that takes as a primitive a cost on experiments.

Posterior-based costs. Let $c : \Delta(\Delta(A)) \rightarrow [0, \infty]$ be a cost on random posteriors such that for all $a \in A$,

$$c(\delta_a) = 0. \tag{30}$$

For choosing an experiment P , the follower incurs a cost $c(B(\alpha, P))$, where α is her conjecture about the seller's strategy.

A posterior-based approach generates a multiplicity of counterintuitive equilibria. To illustrate, let $BR(\alpha) \subseteq B$ be the set of best pure responses to a strategy $\alpha \in \Delta(A)$ of the leader:

$$BR(\alpha) = \arg \max_{b \in B} u_f(\alpha, b).$$

Note that, by (29), for all $a \in A$, $BR(a)$ is a singleton. We denote by $MRB(\alpha)$ the set of mixed best responses:

$$MRB(\alpha) = \{\beta \in \Delta(B) : \text{supp}(\beta) \subseteq BR(\alpha)\}.$$

We also denote by $MRB(\neg a)$ the set of mixed best responses to a strategy of the leader that does not play a :

$$MRB(\neg a) = \bigcup_{\alpha \in \Delta(A \setminus \{a\})} MRB(\alpha).$$

Proposition 12. *Assume (30) for cost of information. For every $a \in A$, if*

$$u_l(a, BR(a)) > \min \left\{ \inf_{\beta \in MRB(\neg a)} \max_{a' \neq a} u_l(a', \beta), \max_{a' \neq a} u_l(a', BR(a)) \right\},$$

then there exists an equilibrium (α, P, σ) such that $\alpha(a) = 1$.

The proposition generalizes Ravid's multiplicity result. In the ultimatum game, a is any offer between 0 and v , and $BR(a)$ is the decision of buying the good. Let b be the decision of not buying the good, which is a best response to the belief that the seller makes an offer above v . We obtain the inequality

$$u_l(a, BR(a)) > \max_{a' \neq a} u_l(a', b).$$

The proof of the proposition is an abstract version of Ravid's arguments.

Proof. Assume first that there is $\beta \in MBR(\neg a)$ such that

$$u_l(a, BR(a)) > \max_{a' \neq a} u_l(a', \beta). \quad (31)$$

Let P be a binary experiment with two possible outcomes, x and y . Assume that $P_a(x) = 1$; for all $a' \neq a$, $P_{a'}(y) = 1$. Choose σ such that $\sigma(BR(a)|x) = 1$; for all $b \in B$, $\sigma(b|y) = \beta(b)$. By (31), a is a best response to (P, σ) . On path, the follower plays the best response to a . In addition, by (30), the cost of information is zero. Thus, (P, σ) is a best response to a . Off path, the follower plays β , which is sequentially rational since $\beta \in MBR(\neg a)$. Overall, (a, P, σ) is an equilibrium, even if we impose sequential rationality.

Assume now that

$$u_l(a, BR(a)) > \max_{a' \neq a} u_l(a', BR(a)). \quad (32)$$

Let P be an uninformative experiment. Choose σ such that for all x , $\sigma(BR(a)|x) = 1$. By (32), a is a best response to (P, σ) . On path, the follower plays the best response to a . In addition, by (30), the cost of information is zero. Thus, (P, σ) is a best response to a . Since P is uninformative, no signal is off-path. Overall, (a, P, σ) is an equilibrium, even if we impose sequential rationality. \square

Experiment-based costs. Let $h : \mathcal{E} \rightarrow [0, \infty]$ be a cost on experiments such

$$h(P) = 0 \text{ if and only if } P \text{ is uninformative.} \quad (33)$$

For choosing an experiment P , the follower incurs a cost $h(P)$, which does not depend on her conjecture about the leader's strategy.

The next proposition, which extends Proposition 8 in the main text, shows that the counterintuitive equilibria that Ravid finds disappear.

Proposition 13. *Assume (33) for cost of information. The following conditions hold:*

(i) *In every equilibrium (α, P, σ) where P is informative,*

$$\sum_a \alpha(a) \int_X u_f(a, \sigma(x)) dP_a(x) > \max_{b \in B} \sum_a \alpha(a) u_f(a, b).$$

In particular, α is not degenerate.

(ii) *For every $a \in A$, there exists an equilibrium (α, P, σ) with $\alpha(a) = 1$ if and only if*

$$u_l(a, BR(a)) > \max_{a' \neq a} u_l(a', BR(a)).$$

Proof. (i). The follower can always obtain the payoff

$$\max_{b \in B} \sum_a \alpha(a) u_f(a, b)$$

by choosing an uninformative experiment, which costs zero. For this deviation not to be profitable, by (33) it must be that

$$\sum_a \alpha(a) \int_X u_f(a, \sigma(x)) dP_a(x) > \max_{b \in B} \sum_a \alpha(a) u_f(a, b).$$

In particular, α cannot be degenerate.

(ii). Let (α, P, σ) be an equilibrium such that $\alpha(a) = 1$. By (i), P is uninformative. Thus, the follower must play the best response to a no matter what signal she observes. For the leader not to have an incentive to deviate, it must be that a is a best response to $b = BR(a)$ —that is,

$$u_l(a, BR(a)) > \max_{a' \neq a} u_l(a', BR(a)).$$

Conversely, take an action a for which the above inequality holds. Let P be an uninformative experiment. Choose σ such that for all x , $\sigma(BR(a)|x) = 1$. By the inequality above, a is a best response to (P, σ) . On path, the follower plays the best response to a . In addition, by (33), the cost of information is zero. Thus, (P, σ) is a best response to a . Since P is uninformative, no signal is off-path. Overall, (a, P, σ) is an equilibrium, even if we impose sequential rationality. \square

References

- Aliprantis, C. D. and K. C. Border**, *Infinite Dimensional Analysis*, 3rd ed., Springer, 2006.
- Bloedel, A. W. and W. Zhong**, “The cost of optimally acquired information,” 2021. Available at <https://sites.google.com/site/alexanderbloedel/home>.
- Cover, T. M. and J. A. Thomas**, *Elements of Information Theory*, 2nd ed., John Wiley & Sons, 2012.
- Denti, T., M. Marinacci, and A. Rustichini**, “The experimental order on random posteriors,” 2022. Available at <https://sites.google.com/site/tommasojdenti/home>.
- , – , and **L. Montrucchio**, “A note on rational inattention and rate distortion theory,” *Decisions in Economics and Finance*, 2020, *43*, 75–89.
- Dillenberger, D., V. Krishna, and P. Sadowski**, “Subjective information choice processes,” 2022. Available at <https://web.sas.upenn.edu/ddill/>.
- Hébert, B. and M. Woodford**, “Rational inattention when decisions take time,” 2021. Available at <https://bhebert.people.stanford.edu>.
- Liese, F. and I. Vajda**, “On divergences and informations in statistics and information theory,” *IEEE Transactions on Information Theory*, 2006, *52*, 4394–4412.
- Matějka, F. and A. McKay**, “Rational inattention to discrete choices: A new foundation for the multinomial logit model,” *American Economic Review*, 2015, *105*, 272–298.
- Morris, S. and P. Strack**, “The Wald problem and the equivalence of sequential sampling and static information costs,” 2019. Available at SSRN 2991567.
- Pomatto, L., P. Strack, and O. Tamuz**, “The cost of information,” 2020. Available at [arXiv:1812.04211](https://arxiv.org/abs/1812.04211).
- Ravid, D.**, “Ultimatum bargaining with rational inattention,” *American Economic Review*, 2020, *110*, 2948–2963.
- Torgersen, E.**, *Comparison of Statistical Experiments*, Cambridge University Press, 1991.