Dynamic information design via deviation rules*

Henrique de Oliveira Rohit Lamba

April 2025

Preliminary

Abstract

This paper studies a general dynamic information design problem and shows how the empirical content therein can be completely characterized using the instrument of deviation rules.

1 Introduction

There is an underlying state of the world distributed according to a fixed prior. An agent observes a signal correlated with this state, takes an action, observes another signal, and takes another action until some finite terminal period. The payoff is realized at the end as a function of the actions taken and the state of the world, which can be non-separable in actions. The problem is described in Figure 1.

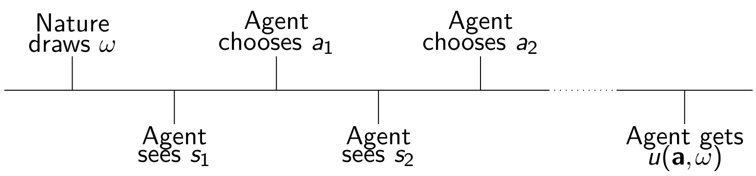

Figure 1: The timeline of signals and actions

The main question we ask is: What distributions over action sequences are feasible, given that the agent is assumed to be Bayesian and expected utility maximizer? Moreover, if we provide agency to the sender of the information through a specific objective function which is different than the agent, what is the optimal dynamic information structure for the sender?

*de Oliveira: Fundação Getúlio Vargas-EESP, henrique.deoliveira@fgv.br; Lamba: Cornell University, rohitlamba@cornell.edu.

This may be christened a *dynamic information design problem* with a single agent. Conceptually, this class of problems explores how information can be used as an instrument to induce rational agents to take actions over time that would be strictly suboptimal if they were exposed only to a static information structure, either at the outset or even at some fixed time in the future. Empirically, they can help characterize the set of distributions observable amongst a population of agents making similar choices over time, without imposing any parametric assumption on the information they are observing. A formal way to describe the empirical content of the Bayesian problem is to identify the restrictions imposed by the joint hypothesis of Bayesian rationality and the specific payoff function u, but without any hypothesis on the information structure.

If $T = 1$, this setup is exactly the Bayesian persuasion problem studied in Kamenica and Gentzkow [2011] or the single-agent version of the information design problem studied in Bergemann and Morris [2016]. Relatedly, Caplin and Martin [2015] put the framework to understanding the empirical content of the static Bayesian decision problem—what joint distributions over actions and states are consistent with a particular information structure; which is extended in Caplin and Dean [2015] to account for Shannon costs of information acquisition.

When the prior is not fixed and is also included as a flexible seed of the arbitrary sequential information structure, de Oliveira and Lamba [2025] characterized the empirical content of the problem for $T > 1$: What class of joint distributions over actions sequences and states of the world can be rationalized for a fixed dynamic decision problem when the analyst does not know the prior or the dynamically arriving signals? It also noted that in many applications at hand, the analyst may not have access to information over realized states of the world, and only be able to record the observed action sequences. To tackle such scenarios, it developed a result for rationalizing the distribution only over action sequences, that is, marginal over actions of the aforementioned joint distribution.

The main contribution of this paper is to further restrict the prior in this analysis. So now the mathematical statement of the problem is to find a marginal distribution over action sequences that, along with the fixed prior, completes a joint distribution that could have been generated by a Bayesian agent that maximizes expected utility. In doing that, it relates to a number of recent studies in both static and dynamic information design.

Analogous questions have been raised in Rehbeck [2023] and Doval et al. [2024], for the case of $T = 1$. These papers also motivate their studies by appealing to the inability of the outside analyst to infer data on realized states, so the analyst must restrict himself to rationalizing distribution over actions. They present distinct duality results, and mapping them to ours for the case of $T = 1$ remains an interesting open question.

In addition, a number of recent papers study the rationalization of optimal stopping times in dynamic models of persuasion; see, for example, Ely and Szydlowski [2020], Hébert and Zhong [2025], Koh and Sanguanmoo [2024], and Saeedi et al. [2025]. These studies, in various contexts and to varying levels of generality, examine the extent to which an agent can be made to wait or take consecutive actions by the principle when their preferences are misaligned.

The novelty we bring to this exciting and burgeoning literature is to evaluate an otherwise complicated set of dynamic obedience constraints through the compact tool of *deviation rules*. To pin down the set of feasible action sequences that can be taken by a rational Bayesian agent, we can proceed in two ways:

1. Guess a binding deviation rule that provides an upper bound on the set of distributions that can be rationalized, and then guess an information structure that achieves that bound. This avoids the cumbersome task of traversing the entire universe of information structures to establish either feasibility or optimality.

2. More systematically, the problem of using deviation rules can be set up as a tractable linear program, and so the binding deviation rule can be computationally determined.

This methodology of studying dynamic information design through deviation rules helps reduce it to a static problem and unify various other models that solve these problems exclusively in the space of information structures. As the reader would notice, there are few assumptions imposed on the dynamic decision problem here; in particular, preferences are allowed to be arbitrary, non-separable over time.

2 Model

2.1 Notation

A *stochastic map* from X to a finite set Y is a function $\alpha : X \to \Delta(Y)$, where $\Delta(Y)$ is the set of probability distributions over Y. We represent the probability assigned to y at the point x by $\alpha(y|x)$. The composition of two stochastic maps $\alpha : X \to \Delta(Y)$ and $\beta : Y \to \Delta(Z)$ is given by

$$\beta \circ \alpha(z|x) = \sum_{y \in Y} \beta(z|y)\alpha(y|x).$$

We can think of a lottery as a stochastic mapping whose domain is a singleton. Therefore, given $\alpha \in \Delta(Y)$ and $\beta : Y \to \Delta(Z)$, we write

$$\beta \circ \alpha(z) = \sum_{y \in Y} \beta(z|y)\alpha(y)$$

to be the probability with which z is chosen by $\beta \circ \alpha$.

For a real-valued function $u : Y \to \mathbb{R}$ and for a lottery $\alpha \in \Delta(Y)$, we denote by $u(\alpha) = \sum_{y \in Y} \alpha(y)u(y)$ the expected value of $u(.)$ under the distribution α.

Throughout the text, we consider a finite number of time periods $t = 1, \ldots, T$. For a collection of sets $(X_t)_{t=1}^T$, we will use the following notation

$$X^t = \prod_{\tau=1}^{t} X_\tau \qquad X = \prod_{\tau=1}^{T} X_\tau$$

with elements $\mathbf{x}^t \in X^t$ and $\mathbf{x} \in X$. Finally, a stochastic map $\alpha : X \to \Delta(Y)$ is said to be *adapted* if the marginal probability of the first t terms of $\mathbf{y}$ depends only on the first t terms of $\mathbf{x}$; formally, it is adapted if the function

$$\sum_{y_{t+1},\ldots,y_T} \alpha(y_1, \ldots, y_t, y_{t+1}, \ldots y_T | x_1, \ldots, x_t, x_{t+1}, \ldots, x_T)$$

is constant in $x_{t+1}, \ldots, x_T$.

2.2 Primitives

The basic model on which the information design problem is layered is based on de Oliveira and Lamba [2025] and is as follows: In each time period t, the agent chooses an action a_t from a finite set A_t. Payoffs are determined after period T by a utility function $u(\mathbf{a}, \omega)$, which depends on the entire action sequence $\mathbf{a} = (a_1, \ldots, a_T) \in A$ and a potentially unknown state of the world ω drawn from a finite set Ω. There are no other restrictions on the utility function.

There is a fixed prior $p \in \Delta(\Omega)$. Thence the agent is informed about the underlying state of the world over time through a sequence of signals. The timeline of the dynamic decision problem is expressed in Figure 1. Every period, before taking an action, the agent observes a signal that is (potentially) correlated with the state of the world and with the signals she has observed in the past. Formally, the sequence of signals is generated by a *sequential information structure*:

Definition 1. *A* **sequential information structure** *is a sequence of finite sets of signals* $(S_t)_{t=1}^{T}$ *and a stochastic mapping* $\pi : \Omega \to \Delta(S)$.[1]

The agent's strategy maps each sequence of signals into a lottery over actions every period, with the restriction that the agent cannot base the choice of an action on signals that have not yet been revealed, which we call adaptedness.

Definition 2. *A* **strategy** *for the agent is an adapted stochastic mapping* $\sigma : S \to \Delta(A)$.[2]

Given the sequential information structure π and agent's strategy σ, the probability that the agent takes a given sequence of actions in each state of the world ω is given by $\sigma \circ \pi(\mathbf{a}|\omega)$. Finally, given a prior $p \in \Delta(\Omega)$, she can evaluate her expected payoff:

$$U(\sigma, \pi, p) = \sum_{\omega \in \Omega} p(\omega) \sum_{\mathbf{a} \in A} \sigma \circ \pi(\mathbf{a}|\omega)\, u(\mathbf{a}, \omega).$$

The agent's problem then is to choose an optimal σ given π and p.

2.3 Feasibility

A key first step in any information design problem is to define the notion of feasibility. In the context of our model this translates to which action sequences and signals can be jointly rationalized by an agent who is best responding to the observed information. The notion of feasibility will thus be described thorough a joint distribution over actions and signals.

Definition 3. *A distribution* $\gamma \in \Delta(A \times \Omega)$ *is said to be* **feasible** *if there exists* (σ, π) *such that*

1. $\sum_{\mathbf{a} \in A} \gamma(\mathbf{a}, \omega) = p(\omega)\ \forall, \omega \in \Omega$,
2. $\sigma \in \arg\max_{\hat{\sigma}} U(\hat{\sigma}, \pi, p)$, *and*
3. $\gamma(\mathbf{a}, \omega) = \sigma \circ \pi(\mathbf{a}|\omega) \cdot p(\omega)\ \forall\, \mathbf{a} \in A, \omega \in \Omega$.

Thus a joint distribution γ is feasible if its marginal over the states of the world is equal to the prior, and there exists a sequential information structure such that, in best responding to it, the agent's optimal strategy generates γ. The empirical content of the dynamic information

[1] We can equivalently define the sequential information structure period-by-period as follows. Let $\pi = (\pi_t)_{t=1}^{T}$ be a family of stochastic mappings where $\pi_1 : \Omega \to \Delta(S_1)$, and $\pi_t : \Omega \times S^{t-1} \to \Delta(S_t)\ \forall\, 2 \le t \le T$. With the exception of zero probability events, which do not affect agent's utlity, we can deduce that the two definitions are equivalent. For a proof, see Lemma 3 in de Oliveira [2018].

[2] As with information structures, an equivalent way to think of the agent's strategy is a family of stochastic mappings $\sigma = (\sigma_t)_{t=1}^{T}$, where $\sigma_1 : S_1 \to \Delta(A_1)$, and $\sigma_t : S^t \times A^{t-1} \to \Delta(A_t)\ \forall\, 2 \le t \le T$. It is possible to deduce one formulation from the other.

design problem then is to characterize the set of feasible joint distributions. To do that we will invoke the notion of a deviation rule.

2.4 Deviation rules

A *deviation rule* is simply an adapted mapping $D : A \to \Delta(A)$, where recollect that being adapted means that the marginal distribution on A^t, the (potentially random) deviation strategy for the first t periods, depends only on A^t, the first t elements of the original strategy from which the agent is deviating. We can think of the deviation rule as a list of alternative actions the agent would take as a function of the actions she originally intended to take. Importantly, a deviation rule is a fully prescribed plan, so that if σ is the original strategy, then $D \circ \sigma(\mathbf{a}|\mathbf{s})$ too is a well-defined strategy.

In what follows, we will first use deviation rules to characterize the empirical content of the dynamic information design problem: It organizes in an intuitive way the system of inequalities that describes the feasible set of joint distributions over signals and actions.

3 Characterization of the empirical content

Fix the set of joint distributions that agree with the prior as the marginal on the states of the world:

$$\Gamma_p = \left\{ \gamma \in \Delta(A \times \Omega) \mid \sum_{\mathbf{a} \in A} \gamma(\mathbf{a}, \omega) = p(\omega)\ \forall\ \omega \in \Omega \right\}.$$

Then, from Definition 3, we know that a joint distribution $\gamma \in \Gamma_p$ is feasible if there exists a sequential information structure π and a strategy σ such that

$$U(\sigma, \pi, p) \geqslant U(\sigma', \pi, p)\ \forall\ \sigma', \text{ and}$$

the joint distribution generated by (σ, π, p) on $A \times \Omega$ is the same as γ_p. Following de Oliveira and Lamba [2025], we can reduce the complexity of the search for feasibility from any sequential information structure to a direct information structure.

Definition 4. *(σ, π) is an* **obedient tuple** *if $S = A$ and $\sigma = Id_A$.*

An obedient tuple is given by an information structure which recommends an action, and a strategy of the agent which always obeys the recommendation. When an action sequence can be deemed feasible with an obedient tuple, we say that it is *obediently feasible*. The obedience principle for dynamic information design then immediately follows.

Lemma 1 (Obedience principle). *If $\gamma_p \in \Gamma_p$ is feasible, then it is obediently feasible.*

Using the obedience principle, and the concept of deviation rule, the empirical content of the model can be readily characterized as follows.

Theorem 1. *A distribution $\gamma_p \in \Gamma_p$ is feasible if and only if there exists a deviation rule $D : A \to \Delta(A)$ such that*

$$\sum_{\mathbf{a},\omega} \left[u(\mathbf{a},\omega) - u(D(\mathbf{a}),\omega)\right] \gamma_p(\mathbf{a},\omega) \geqslant 0.$$

A simple intuition follows: If γ_p is feasible, then from Lemma 1, it is obediently feasible. Under such an obedient feasibility, since $S = A$, any alternative strategy $\hat{\sigma}$ is an adapted mapping from A to A, and hence a deviation rule, and the obedience constraint reads as stated in the result above. If, conversely, γ is not feasible, then the candidate obedient strategy is not optimal, and there must be an alternative strategy—a deviation rule—that improves upon it, that is satisfies the opposite of the inequality in the result. This proves the dual characterization stated in Theorem 1.

The result is a direct dynamic generalisation of single-player obedience constraints in information design (see surveys by Bergemann and Morris [2019] and Kamenica [2019]), restricted here to single agent problem. In the lexicon of that literature, all distributions $\gamma \in \Delta(A \times \Omega)$ that satisfy the above inequality for all deviation rules can be supported as a "Bayes correlated equilibrium" of our decision problem.[3] This static version of the constraint has also been used by Caplin and Martin [2015] to characterize the empirical content of Bayesian decision problem faced by a large population of agents. As we discuss in our earlier work, a similar interpretation to the dynamic problem as well in characterizing the empirical content of arbitrary dynamic decision problems with varying sequential information structures.

4 Feasible actions

The main objective of this paper is to characterize the feasible marginals on actions sequences for any fixed prior. As we argued in the introduction, in many realistic scenarios, it may not be feasible for an analyst or an econmetrician to jointly observe the action sequence taken by an agent and the realization of the corresponding state of the world, even though the prior on the states is well known. To that end, the analyst must have access to a technique that allows her

[3] Just as the objective of the linear program in information design is to identify the set of binding obedience constraints, the objective of the above result is to identify the critical deviation rule.

to characterize the empirical content of the model with access to the prior and the observable action sequences. We start by updating the notion of feasibility to adjust to this reality.

Definition 5. *A distribution* $\alpha \in \Delta(A)$ *is* **feasible** *if there exists* (σ, π) *such that*

1. $\sigma \in \arg\max_{\hat{\sigma}} U(\hat{\sigma}, \pi, p)$, *and*
2. $\sigma \circ \pi \circ p = \bar{\gamma}$.[4]

Now, define the set of joint distributions with fixed marginals:

$$\Gamma_{\alpha,p} = \left\{ \gamma \in \Delta(A \times \Omega) \mid \sum_{\mathbf{a}\omega\in\Omega} \gamma(\mathbf{a}, \omega) = \alpha(\mathbf{a}) \ \forall\ \mathbf{a} \in A, \& \sum_{\mathbf{a}\in A} \gamma(\mathbf{a}, \omega) = p(\omega) \ \forall\ \omega \in \Omega \right\}.$$

The main result, a tight characterization of the set of feasible joint distributions that agree with given marginals, follows.

Theorem 2. *A distribution* α *over action sequences* $\alpha \in \Delta(A)$ *is feasible under a decision problem* $u : A \times \Omega \to \mathbb{R}$ *with prior* $p \in \Delta(\Omega)$ *if and only if there exists a* $\varphi : \Omega \to \mathbb{R}$ *such that* $\sum_{\omega} p(\omega)\varphi(\omega) = 0$ *and*

$$\sum_{a} \min_{\omega} \left[u(\mathbf{a}, \omega) - u(D(\mathbf{a}), \omega) + \varphi(\omega)\right] \alpha(\mathbf{a}) \geqslant 0.$$

Proof. Let

$$\Gamma = \left\{ \gamma \in \Delta(\Omega \times A) \mid \gamma(\omega) = p(\omega),\ \gamma(a) = \alpha(a) \right\},$$

be the set of joint distribution consistent with both marginals p and α. By Lemma 1, the question can be restated as: "is there a $\gamma \in \Gamma$ such that

$$\sum_{\omega,a} \left[u(a, \omega) - u(D(a), \omega)\right] \gamma(\omega, a) \geqslant 0$$

for all deviation rules D?" In other words, is it true that

$$\max_{\gamma\in\Gamma} \min_{D} \sum_{\omega,a} \left[u(a, \omega) - u(D(a), \omega)\right] \gamma(\omega, a) \geqslant 0?$$

Using the minimax theorem (see Sion [1958]) this question is equivalent to asking whether

$$\min_{D} \max_{\gamma\in\Gamma} \sum_{\omega,a} \left[u(a, \omega) - u(D(a), \omega)\right] \gamma(\omega, a) \geqslant 0.$$

[4] Since $\sigma \circ \pi \circ p$ generates a joint distribution over $A \times \Omega$, this essentially means that the marginal of $\sigma \circ \pi \circ p$ on A equals $\bar{\gamma}$.

The maximization problem inside is an optimal transport problem and we can calculate its dual to obtain

$$\max_{\gamma\in\Gamma}\sum_{\omega,a}\left[u\left(a,\omega\right)-u\left(D\left(a\right),\omega\right)\right]\gamma\left(\omega,a\right)=\min_{(\varphi,\psi)\in C}\sum_{\omega}\varphi\left(\omega\right)p\left(\omega\right)+\sum_{a}\psi\left(a\right)\alpha\left(a\right)$$

where $C=\{(\varphi,\psi)\,|\forall\omega,a,\ \varphi\left(\omega\right)+\psi\left(a\right)\geqslant u\left(a,\omega\right)-u\left(D\left(a\right),\omega\right)\}$.

For a given φ, the ψ that solves the minimization above is

$$\psi\left(a\right)=\max_{\omega}\left[u\left(a,\omega\right)-u\left(D\left(a\right),\omega\right)-\varphi\left(\omega\right)\right].$$

Thus, the question becomes: is it true that for every φ and D, we have

$$\sum_{a}\max_{\omega}\left[u\left(a,\omega\right)-u\left(D\left(a\right),\omega\right)-\varphi\left(\omega\right)\right]\alpha\left(a\right)+\sum_{\omega}\varphi\left(\omega\right)p\left(\omega\right)\geqslant 0$$

The negation of this statement (meaning that α cannot be rationalized) asks: is there a φ and D such that

$$\sum_{a}\min_{\omega}\left[u\left(D\left(a\right),\omega\right)-u\left(a,\omega\right)+\varphi\left(\omega\right)\right]\alpha\left(a\right)-\sum_{\omega'}\varphi\left(\omega'\right)p\left(\omega'\right)>0$$

Moreover, since adding a constant to φ does not alter the inequality, we can normalize it and restate the question as: is there a D and a φ such that $\sum_{\omega}\varphi\left(\omega\right)p\left(\omega\right)=0$ and

$$\sum_{a}\min_{\omega}\left[u\left(D\left(a\right),\omega\right)-u\left(a,\omega\right)+\varphi\left(\omega\right)\right]\alpha\left(a\right)>0.$$

□

Notice that, taking $\varphi=0$, the condition becomes the same as when we let the prior be arbitrary (see de Oliveira and Lamba [2025]). Restricting the prior to p then makes it easier to find a deviation rule that rejects the model, because of the flexibility afforded by the choice of φ. We can think of φ as a transfer between states, with the exchange rate given by p.

Now Theorem 2 presents a characterization of the set feasible action profiles in an arbitrary information design problem. In many settings, the modeler is interested in not just interested in characterizing the entire set of feasible distributions, rather by equipping the "sender" or the "principal" with some preferences, she wants to find out the optimal information structure too. To show the workings of such an exercise, we now present a example that seeks to maximize the probability of an apparently dominated action sequence—that is one which the sender can never induce the agent to take with a static information structure, but can with some positive probability, when she has access to a dynamic information structure.

5 A two-pronged approach

Theorem 2 gives an if and only if result, so if we can compute the function

$$\Delta(\alpha) = \max_{D,\varphi} \sum_{\mathbf{a}} \min_{\omega} \left[u\left(D\left(\mathbf{a}\right), \omega\right) - u\left(\mathbf{a}, \omega\right) + \varphi\left(\omega\right) \right] \alpha\left(\mathbf{a}\right),$$

we know precisely the set of α's that can be induced by a suitable choice of information structure. A persuasion problem is simply a maximization problem over that set of α's.

In practice, however, just as the set of information structures can be large, the set of all deviation rules can also be large, so computing $\Delta(\alpha)$ can be tiresome. A more practical method is to combine both approaches. First use the idea of the theorem to find an upper bound on $\Delta(\alpha)$. If for some deviation rule D, we have

$$\sum_{\mathbf{a}} \min_{\omega} \left[u\left(D\left(\mathbf{a}\right), \omega\right) - u\left(\mathbf{a}, \omega\right) + \varphi\left(\omega\right) \right] \alpha\left(\mathbf{a}\right) > 0,$$

then α cannot be feasible. This means that, by considering a subset of deviation rules, we can rule out many distributions.

Second, to show that α can be rationalized, it is sufficient to find a joint distribution $\gamma \in \Gamma_{\alpha,p}$ that satisfies all obedience constraints. If we can directly show that everything that was not ruled out by our subset of deviation rules can be rationalized, we have a complete characterization of the set of feasible distributions over actions.

6 An example of information design through deviation rules

In this section, we illustrate through an example how our result can be used to solve dynamic persuasion (or information design) problems. We consider the following two period decision problem based on Ely and Szydlowski [2020]: The agent faces a task and has a choice of working or not working in each period, where working has an effort cost of c. The agent is unsure about whether the task is easy or hard, so $\Omega = \{E, H\}$ with prior $p = (p_E, p_H)$. If it is easy, a single round of work suffices for the agent to get the reward R; if it is hard, the agent will only get the reward R if they work in both periods. Payoffs and decision tree are represented as follows:

	n	(w,n)	(w,w)
hard	0	$-c$	$R-2c$
easy	0	$R-c$	$R-2c$

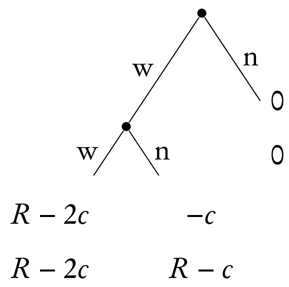

The question we ask is: what is the maximum probability with which we can induce the agent to work twice? So in the language of information design, the "sender's" preference is simply to maximize the probability of the actions sequence (w, w). To avoid uninteresting cases, we assume $0 < c < R < 2c$ and $0 < p_E, p_H < 1$.

It is easy to that the action sequence (w, w) is apparently dominated—there exists an alternate action sequence, simply choosing n in the first period, such that the agent gets a strictly higher payoff in each state of the world. Therefore, no static information structure can deem the selection of this action sequence feasible, but a dynamic information structure can. Moreover, since (w, w) is apparently dominated, no (dynamic) information structure for any interior prior can induce the agent to pick it with probability 1. So the information design exercise is to pin down the highest feasible probability of an optimizing agent selecting (w, w) as a function of p.

To simplify notation, we let $\kappa = \frac{c}{R}$ in what follows.

6.1 Deviating from work to not work in the first period

First we consider the first-period deviation rule, from work to not work. In this case, with a fixed prior, this deviation rule rejects a distribution over action sequences if there exists a φ_E, φ_H such that $\varphi_E p_E + \varphi_H p_H = 0$ and

$$\min\{\varphi_E, \varphi_H\}\,\gamma(n) + \min\{2c - R + \varphi_E, 2c - R + \varphi_H\}\,\gamma(ww) + \min\{c - R + \varphi_E, c + \varphi_H\}\,\gamma(wn) > 0.$$

Letting $\varphi_E = \frac{k}{p_E}$ and $\varphi_H = -\frac{k}{p_H}$, we get that the deviation rule does *not* reject a distribution

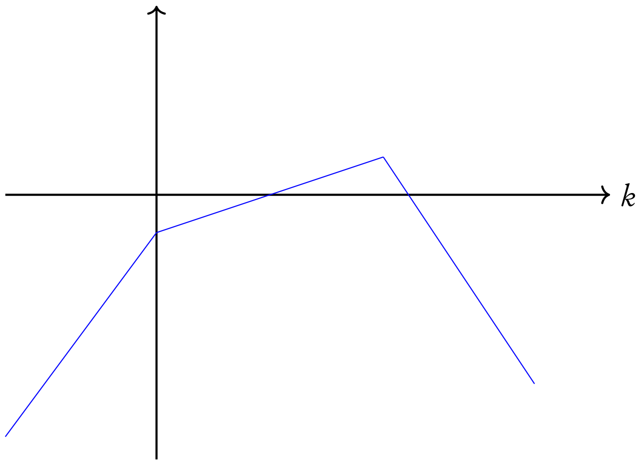

Figure 2: The function is negative if and only if it is negative at both kinks.

over action sequences if, for all $k \in \mathbb{R}$, we have

$$\min\left\{\frac{k}{p_E}, -\frac{k}{p_H}\right\}\gamma(n) + \left(2c - R + \min\left\{\frac{k}{p_E}, -\frac{k}{p_H}\right\}\right)\gamma(ww) + \min\left\{c - R + \frac{k}{p_E}, c - \frac{k}{p_H}\right\}\gamma(wn) \leqslant 0.$$

The left-hand-side is a piece-wise linear function that is increasing for low values of k, decreasing for high values of k, and has two kinks, at $k = 0$ and $k = Rp_E p_H$. Thus, to prove that this function is always negative, it is enough to show that its values at the two kinks are negative. At $k = 0$ this gives us the condition

$$(2c - R)\gamma(ww) + (c - R)\gamma(wn) \leqslant 0.$$

or

$$\gamma(ww) \leqslant \frac{(R - c)}{(2c - R)}\gamma(wn) \tag{1}$$

This shows that an upper bound for the value of $\gamma(ww)$ is given by

$$\gamma(ww) \leqslant \frac{(R - c)}{(2c - R)}(1 - \gamma(ww))$$

which becomes

$$\gamma(ww) \leqslant \frac{R - c}{c} = \frac{1}{\kappa} - 1. \tag{2}$$

At $k = Rp_E p_H > 0$, we get the condition

$$-\frac{Rp_E p_H}{p_H}\gamma(n) + \left(2c - R - \frac{Rp_E p_H}{p_H}\right)\gamma(ww) + \left(c - \frac{Rp_E p_H}{p_H}\right)\gamma(wn) \leqslant 0$$

or

$$\left(2c - R - Rp_E\right)\gamma(ww) \leqslant Rp_E\gamma(n) + \left(Rp_E - c\right)\gamma(wn).$$

Substituting $\gamma(n) = 1 - \gamma(ww) - \gamma(wn)$, we can write this inequality purely in terms of $\gamma(ww)$ and $\gamma(wn)$:

$$(2c - R)\gamma(ww) \leqslant Rp_E - c\gamma(wn).$$

Using (1), we get

$$(2c - R)\gamma(ww) \leqslant Rp_E - c\frac{2c - R}{R - c}\gamma(ww)$$

or

$$\gamma(ww) \leqslant \frac{R - c}{2c - R}p_E = \frac{1 - \kappa}{2\kappa - 1}(1 - p_H). \tag{3}$$

6.2 Deviating from Work to not work in the second period

Now we consider another deviation rule, that deviates from work to not work in the second period, and does not deviate from the other sequences. This deviation rejects a distribution over action sequences if there exists a φ_E, φ_H such that $\varphi_E p_E + \varphi_H p_H = 0$ and

$$\min\{\varphi_E, \varphi_H\}[\gamma(n) + \gamma(wn)] + \min\{c + \varphi_E, c - R + \varphi_H\}\gamma(ww) > 0.$$

Letting $\varphi_E = -\frac{k}{p_E}$ and $\varphi_H = \frac{k}{p_H}$, we get that the deviation rule does *not* reject a distribution over action sequences if, for all $k \in \mathbb{R}$, we have

$$\min\left\{-\frac{k}{p_E}, \frac{k}{p_H}\right\}[\gamma(n) + \gamma(wn)] + \min\left\{c - \frac{k}{p_E}, c - R + \frac{k}{p_H}\right\}\gamma(ww) \leqslant 0.$$

As before, this is a piecewise linear concave function with two kinks, at $k = 0$ and at $k = Rp_E p_H$. At $k = 0$ we get the inequality

$$(c - R)\gamma(ww) \leqslant 0$$

which holds trivially. At $k = Rp_E p_H > 0$ we get

$$-Rp_H[\gamma(n) + \gamma(wn)] + \left(c - Rp_H\right)\gamma(ww) \leqslant 0.$$

and substituting $\gamma(n) + \gamma(wn) = 1 - \gamma(ww)$, we get

$$\gamma(ww) \leqslant \frac{R}{c} p_H \tag{4}$$

Putting together the inequalities (2),(3), and (4), we get the following bound

$$\gamma(ww) \leqslant \begin{cases} \frac{p_H}{\kappa} & \text{for } p_H \leqslant 1-\kappa \\ \frac{1}{\kappa} - 1 & \text{for } 1-\kappa < p_H \leqslant \frac{1}{\kappa} - 1 \\ \frac{1-\kappa}{2\kappa-1}\left(1-p_H\right) & \text{for } \frac{1}{\kappa} - 1 < p_H \end{cases}$$

6.3 Lower bound by construction of information structure

So far, we have obtained an upper bound on the probability of the agent choosing to work twice. To show that this upper bound is tight, we now construct an information structure that reaches it. We divide the construction in cases, depending on how the prior belief relates to κ.

6.3.1 $p_H \leqslant 1-\kappa$

In this case, the agent believes that the problem is easy with a high probability, so the agent already has a high incentive to work in the first period. Therefore let's start assuming that the agent gets no information in the first period.

In the second period, assume that the agent is always told to work when the problem is hard and let a be the probability that the agent is told to work and the problem is easy. To achieve the upper bound, we must have $a + p_H = \frac{p_H}{\kappa}$, or $a = p_H\left(\frac{1}{\kappa} - 1\right)$. We can check that the obedience constraint is indeed satisfied: the probability that the problem is hard when the agent is told to work is $\frac{p_H}{a+p} = \kappa$, making the agent exactly indifferent between working and not working when told to work. This gives the agent an ex-ante expected utility of

$$(a + p_H)(R - 2c) + (1 - a - p_H)(R - c) = R - \left(1 + \frac{p_H}{\kappa}\right)c.$$

which is positive so long as $p_H \leqslant 1-\kappa$, so the agent also has an incentive to work in the first period.

6.3.2 $1-\kappa < p_H \leqslant \frac{1}{\kappa}-1$

When $p_H > 1-\kappa$, the previous information structure would not give the agent enough incentive to work in the first period. We can take care of this by changing a, the probability that the agent is told to work twice and the state is easy. To achieve the upper bound, we write $a+p_H = \frac{1}{\kappa}-1$. The agent's belief that the state is hard when told to work the second time is

$$\frac{p_H}{a+p_H} = \frac{p_H}{\frac{1}{\kappa}-1} > \kappa,$$

where the inequality follows from $p_H > 1-\kappa$. This means that the agent has indeed an incentive to work when told to do so.

Notice that when $p_H = \frac{1}{\kappa}-1$, $a=0$, so the agent is given full information.

6.3.3 $p_H > \frac{1}{\kappa}-1$

When the agent believes the problem is hard with high probability, they will not have an incentive to work in the first period, unless they receive some information in the first period. Let $b \leqslant p_H$ be the probability that the agent is told to work and the problem is hard in the first period. We choose b so that, after being told to work in the first period, the agent believes the probability of the state being hard is

$$\frac{b}{b+1-p_H} = \frac{1}{\kappa}-1.$$

In the second period, we proceed as in the previous case: the agent is given full information, getting an expected payoff of zero. Thus the probability that the agent ultimately chooses to work twice is given by

$$b = (1-p_H)\frac{(\frac{1}{\kappa}-1)}{2-\frac{1}{\kappa}} = \frac{1-\kappa}{2\kappa-1}(1-p_H)$$

which is precisely the upper bound we found before. This finishes the proof that the bound is tight.

7 Final remarks

This paper presents a duality result that completely characterizes the empirical (informational) content of a general dynamic decision problem. It argues that the instrument of deviation rule can help unify disparate studies that explore similar questions in the burgeoning literature on static and dynamic information design.

References

D. Bergemann and S. Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.

D. Bergemann and S. Morris. Information design: a unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.

A. Caplin and M. Dean. Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203, 2015.

A. Caplin and D. Martin. A testable theory of imperfect perception. *The Economic Journal*, 125:184–202, 2015.

H. de Oliveira. Blackwell's informativeness theorem using diagrams. *Games and Economic Behavior*, 109:126–131, 2018.

H. de Oliveira and R. Lamba. Rationalizing dynamic choices. Fundação Getúlio Vargas-EESP and Cornell University, 2025.

L. Doval, R. Eilat, T. Liu, and Y. Zhou. Revealed information. Columbia University and Ben Gurion University, 2024.

J. C. Ely and M. Szydlowski. Moving the goalposts. *Journal of Political Economy*, 128(2):468—506, 2020.

B. Hébert and W. Zhong. Engagement maximization. Stanford University, 2025.

E. Kamenica. Bayesian persuasion and information design. *Annual Review of Economics*, 11: 249–272, 2019.

E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

A. Koh and S. Sanguanmoo. Attention capture. MIT, 2024.

J. Rehbeck. Revealed bayesian expected utility with limited data. *Journal of Economic Behavior & Organization*, 207:81–95, 2023.

M. Saeedi, Y. Shen, and A. Shourideh. Getting the agent to wait˙ Carnegie Mellon University, 2025.

M. Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.