Apparent Mistakes*

Henrique de Oliveira   Rohit Lamba

May 2017

Preliminary draft

Abstract

An analyst wants to predict which sequence of actions an agent will take over time. He knows the set from the agent's payoff is drawn, but does not know how informed the agent is about her payoffs. Assuming that the agent acts according to a fixed belief about her payoff function, the analyst would predict that no sequence of actions which is strictly dominated will be taken. But if the agent gradually learns her utility function over time, it is possible that she chooses a sequence of actions which is strictly dominated. Only truly dominated sequences of actions- those that are never optimal under any gradual learning- can be ruled out. Thus observing a strictly dominated sequence of actions may appear to be a mistake while being plausibly explained by gradual learning.

We show that a sequence of actions is truly dominated precisely when the analyst can rule it out by designing a deviation plan that is measurable with respect to the agent's actions and does strictly better.

1 Introduction

An agent takes a sequence of actions over finite time: $\mathbf{a} = (a_t)_{t=1}^T$. She obtains a terminal payoff as a function of the actions taken and an underlying state of the world; the state is unknown to the agent. An analyst observes the sequence of actions. He knows the payoff function $u : A \times \Omega \to \mathbb{R}$, where $A = \prod_{t=1}^{T} A_t$ is the set of actions and Ω is the set of possible states. However, he is not aware of what the agent knows about her payoffs; more specifically he does not know how informed the agent is about ω.

Suppose the analyst uses the standard model that fixes the agent's belief about the state of the world before any action is taken. The agent can have a prior and receive signals before taking a sequence of actions; we call this the *fixed belief* model. It will be easy to see that under such reasoning, the agent will never choose a *strictly dominated* action sequence.

A sequence of actions $\mathbf{a}$ is *strictly dominated* if there exists another sequence of actions $\mathbf{b}$ that provide a strictly higher payoff in every state of the world: $u(\mathbf{b}, \omega) > u(\mathbf{a}, \omega)\ \forall\ \omega \in \Omega$.[1] If the

*We are indebted to Vijay Krishna and Ran Shorrer for critical comments that helped formulate the idea.

[1] Technically speaking we allow for any randomized strategy, so $\mathbf{b} \in \Delta(A)$.

analyst were to use a fixed belief model, he would brand a strictly dominated sequence of actions as a *mistake*.

Our analyst's model allows for the possibility of *gradual learning*. Every period before taking an action, the agent observes a signal that is correlated with the underlying state of the world and with the signals she has observed in the past. These signals are exogenous in that they are independent of the actions taken by the agent. The class of all such signals will be modeled through a *sequential information structure*. The agent is a Bayesian expected utility maximizer. Under this formulation the agent may take a strictly dominated action with positive probability. Consider the following example.

Example 1. *A firm (agent) can invest* (I) *in a risky technology or go safe* (S). *The risky technology pays off in the good state* (G) *but hurts in the bad* (B). *If the firm chooses* I, *it gets an opportunity to either pull back* (b) *or go ahead* (a) *just before the investment is completely locked in. Pulling back is costly but it prevents the firm from incurring the larger negative shock of a bad investment. The firm can receive a signal about the state of the world before deciding on* I *or* S *and then again before deciding on* b *or* a. *A parametrization of payoffs is expressed in Table 1.*

	Ib	Ia	S
G	0	3	1
B	0	-3	1

Table 1: Action sequence Ib is strictly dominated but can be taken with a positive probability

The agent may receive a signal that informs her that state G is highly likely, which makes action I optimal in the first period. But, the second signal may bring down the likelihood of G making the agent take action b *in the second period. Note that* Ib *is strictly dominated by* S. *So, if we were to allow for the fixed belief model only, the firm would never choose* Ib. *Allowing for a sequential information structure opens up the possibility of a strictly dominated action being rationally taken with a positive probability.*

What actions can the analyst completely rule out under the assumption of a Bayesian maximizer? If a sequence of actions is such that it can never be taken by the agent with positive probability under any sequential information structure, we say that it is *truly dominated*. Such actions cannot be rationalized by the analyst. Interestingly no action sequence in Example 1 is truly dominated. Consider the following second example.

Example 2. *A firm can choose to invest today in one of the two technologies* x *or* y; *the former generates great rewards in state* X *and the latter generates great rewards in state* Y. *The firm can also wait* (w) *to acquire more information before eventually choosing between the two technologies. However, it uniformly looses out to some competitor in case of waiting and receives a fraction* δ *of the rewards it would have, had it chosen to invest right away. A parametrization of payoffs is expressed in Table 2.*

The firm may receive a poor signal in the first period- say both states are equally likely. If the firm expects to get a precise signal in the second period, then for high enough value of δ *it will decide to wait,*

	x_1	y_1	wx_2	wy_2
X	3	1	3δ	δ
Y	1	3	δ	3δ

Table 2: Action sequences wx_2 and wy_2 are strictly dominated for all $0 < \delta < 1$. They are truly dominated iff $\delta < 2/3$

even though wx_2 is strictly dominated by x_1, and wy_2 by y_1. Importantly though, waiting can only be rationalized till a threshold value of δ. We can show that for all $\delta < 2/3$, action sequences wx_2 and wy_2 are truly dominated. Thus, for $2/3 \leq \delta < 1$, an analyst with a fixed belief model would regard waiting to be a mistake, whereas the gradual learning model would be able to rationalize the choice.

What is the set of truly dominated action sequences? Trudging through the class of all possible sequential information structures may be a hard exercise. Is there an equivalent way to determine whether a sequence of actions is truly dominated without having to appeal to all possible signals that the agent can observe? Example 2 offers a hint. Note that for $0 < \delta < 2/3$, the *mixed strategy* $\sigma = \left(\frac{1}{2}x_1, \frac{1}{2}y_1\right)$ dominates w in the first period- if for *any first period signal* the agent chose w, she can do strictly better by switching to σ. Thus, the choice of w in the first period is truly dominated for any action sequence.

There are two challenges in generalizing the aforementioned argument. First, the contingent deviation recommended to the agent only had to suggest an alternative first period action. How does one define such deviations in more complex situations? And second, is the construction of such contingent deviations also necessary for true dominance? Let us consider a third example.

Example 3. *Consider the payoff function described in Table 3. Here $\Omega = \{\omega_1, \omega_2\}$, $A_1 = \{a, b, c\}$, and $A_2 = \{l, r\}$.*

	al	ar	bl	br	cl	cr
ω_1	4	-4	8	2	-9	2
ω_2	-4	4	-9	2	8	2

Table 3: Action sequences al and ar are truly dominated

The action sequences al and ar are truly dominated. Indeed, if the agent chooses either of those sequences with positive probability, then there is a deviation which would give that agent strictly higher utility. To see this, consider the following deviation plan:

1. *whenever the agent would have chosen action a in the first period, instead randomize with equal probability between b and c.*
2. *in the second period,*
   (a) *whenever the agent would have chosen al, choose l if the result of the first-period randomization was b and choose r if it was c;*
   (b) *whenever the agent would have chosen ar, choose r if the result of the first-period randomization was b and choose l if it was c.*

Deviating in this fashion is clearly feasible for the agent, since it depends on information (or signals) that the agent already possess. Moreover, this deviation guarantees an increase in payoff:

1. *whenever the agent would have chosen al, the expected payoff after the deviation is*

$$\frac{1}{2}(8,-9)+\frac{1}{2}(2,2)\gg(4,-4)$$

2. *whenever the agent would have chosen ar, the expected payoff after the deviation is*

$$\frac{1}{2}(2,2)+\frac{1}{2}(-9,8)\gg(-4,4)\,.$$

In general, if there exists a deviation plan that strictly improves the payoff for a certain sequence of actions, we can use an argument analogous to the one above to conclude that this sequence of actions is truly dominated. Our main result states that this characterization is tight in that it is not only sufficient but also necessary for true dominance.

Theorem. A sequence of actions is truly dominated if and only if there exists a deviation plan that improves on that sequence of actions.

The above statement is somewhat informal for the term *deviation plan* has to be formally defined which we do so in Section 3. In a nutshell, the deviation plan has to be adapted to the filtration generated by the original strategy from which the agent is deviating. As we saw in Example 3, the deviation was constructed so that agent is prescribed a different action in the first period, and then as a function of his original action sequence and first period deviation, she is prescribed a second period deviation.

The deviation plan may involve randomization over multiple periods. But, importantly, it only uses as information the actions that the agent would otherwise have taken. Therefore this characterization does not require any direct consideration of possible information structures. In the case of a single time period, this characterization reduces to the classical result that an action is never a best response if and only if it is dominated by some mixed strategy.[2] Hence, if a mixed strategy involves playing a dominated action with positive probability, it can be improved by a deviation plan that switches the action to the mixed strategy that dominates it and leaves all other actions the same.

The model allows us to say a bit more. If we are willing to fix a prior, the analyst can precisely pin down an upperbound on the probability with which a strictly dominated action can be taken for any information structure. It amounts to solving for a dynamic version of the persuasion problem àla Kamenica and Gentzkow [2011] where the "sender" wants to maximize the probability of taking the strictly dominated action. If we thought of the analyst as someone observing a large population making the dynamic decision problem, this result tells us the maximum possible

[2]See Lemma 60.1 in Osborne and Rubinstein [1994]

fraction of people that could plausible make an apparent mistake, that is take a strictly but not truly dominated action.

There is also a message on behavioral dynamic choice. The set of *apparent mistakes* may be construed to be "irrational" or "suboptimal" under a fixed belief model. The gradual learning model provides a novel theory that can potentially shed light on the choice of seemingly dominated actions in situations where time plays an important role. Why do people pay upfront for an expensive gym membership and not utilize it sufficiently to justify that choice (DellaVigna and Malmendier [2006])? It could be because they expected to have more time on their hands, or care more about their health status than they eventually end up having or doing. One can think of similar motivations for usage of credit cards with high interests on outstanding balances (Ausubel [1991]).

Related literature. Kamenica and Gentzkow [2011], Bergemann and Morris [2016], Ely [2017], Bergemann and Morris [2017] [to be completed]

2 Model and definitions

An agent takes a sequence of actions: $a_t \in A_t$, at each point in time $t \leq T$. The set of possible actions is finite: $|A_t| < n$ for some $n \in \mathbb{N}$. Let $A^t = \prod_{\tau=1}^{t} A_t$ be the set of all action sequences up to time t, and $A = A^T$ be the meta set of actions. Also, denote a representative element of A by by $\mathbf{a} = (a_t)_{t=1}^T$. There is a perfectly persistent underlying state of the world $\omega \in \Omega$, unknown to the agent. For an action sequence $\mathbf{a}$ and state ω, a terminal payoff $u(\mathbf{a}, \omega)$ is realized at the end of period T. Similarly, payoff from a randomized action sequence $\alpha \in \Delta(A)$ is given by $u(\sigma, \omega) = \sum_{\mathbf{a} \in A} \sigma(\mathbf{a}) u(\mathbf{a}, \omega)$, where $\sigma(\mathbf{a})$ refers to the probability of action sequence $\mathbf{a}$ under α.

Definition 1. *An action sequence* $\mathbf{a} \in A$ *is* **strictly dominated** *if there exists a randomized action sequence* $\sigma \in \Delta(A)$ *such that*

$$u(\sigma, \omega) > u(\mathbf{a}, \omega) \;\forall\; \omega \in \Omega$$

Dominance thus defined provides a clear reasoning for the suboptimality of the dominated action. An analyst with a fixed belief model would rule out all strictly dominated actions as being implausible. To that end, we set the stage for sequential information design.

Definition 2. *A* **sequential information structure** *is defined by* (π, S), *where* $S = \prod_{t=1}^{T} S_t$, *is a finite set of signals, and* $\pi : \Omega \to \Delta(S)$ *is a signal distribution.*

To help digest the rather sparse exposition in Definition 2, consider the following equivalent formulation. Let $S^t = \prod_{\tau=1}^{t} S_\tau$ with $S = S^T$ as the meta set of sequentially arriving signals. A sequential information structure is given by (π, S), where $\pi = (\pi_t)_{t=1}^T$ is a family of signal distributions such that $\pi_1 : \Omega \to \Delta(S_1)$, and $\pi_t : \Omega \times S^{t-1} \to \Delta(S_t) \;\forall\; 2 \leq t \leq T$. Taking the set of signals S as given, we will abuse notation slightly and refer to π itself as the information structure.

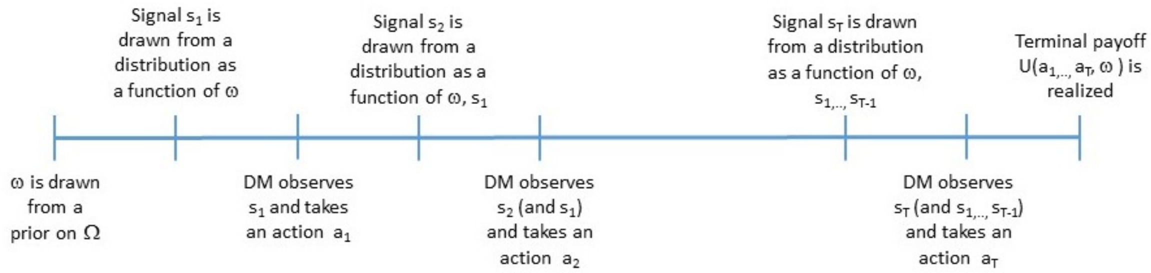

Figure 1: The timeline of signals and actions

Our formulation of a sequential information structure can be thought of as a dynamic generalization of a (static) persuasion problem àla Kamenica and Gentzkow [2011] or (static) information design problem àla Bergemann and Morris [2016].

Every period before taking an action, the agent observes a signal that is correlated with the underlying state of the world and with the signals she has observed in the past. Note that these signals are exogenous in that they are independent of the actions taken by the agent. The timeline of the dynamic decision problem is expressed in Figure 1.

We say that a function $f : S \to \Delta(A)$ is adapted if the marginal distribution on A^t depends only on S^t. Formally, the agent's strategy can be defined as follows.

Definition 3. *A strategy for the agent is an adapted mapping $\alpha : S \to \Delta(A)$.*

As with information structures, an equivalent way to think the agent's strategy is a collection of mappings $\alpha = (\alpha_t)_{t=1}^{T}$, where $\alpha_1 : S_1 \to \Delta(A_1)$, and $\alpha_t : S^t \times A^{t-1} \to \Delta(A_t)\ \forall\ 2 \leq t \leq T$.

Next, we can define a composition mapping: $\alpha \circ \pi : \Omega \to \Delta(A)$, where we write $\alpha \circ \pi(\mathbf{a}|\omega)$ to denote the probability with which action sequence $\mathbf{a}$ is picked by the agent under strategy α when faced with a sequential information structure (π, S) and underlying state ω.

Finally, equip the agent with a prior $p \in \Delta(\Omega)$; so we can evaluate her expected payoff:.

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

Her decision problem boils down to choosing the optimal strategy α^*, given a sequential information structure (π, S) and prior p in order to maximize her expected payoff $U(\alpha, \pi, p)$. We say that an action sequence is truly dominated if it cannot be chosen with a positive probability by an optimizing agent for any possible information structure; defined formally thus.

Definition 4. *An action sequence* $\mathbf{a} \in A$ *is* **truly dominated** *if for any prior p, any information structure (π, S) and any agent's strategy $\alpha : S \to \Delta(A)$ such that $\sum_{\omega\in\Omega} \alpha \circ \pi(\mathbf{a}|\omega) p(\omega) > 0$, there exists another strategy $\beta : S \to \Delta(A)$ such that*

$$\sum_{\omega\in\Omega} \beta \circ \pi(\mathbf{a}|\omega) p(\omega) = 0 \quad \textit{and} \quad U(\beta, \pi, p) > U(\alpha, \pi, p)$$

There is an implicit assumption in our definition of true dominance. We assume that the prior is interior: $p \in \Delta_+(A)$. This for the following reason. We assume that all states of the world which are *possible*, those which the true data generating process weighs with positive probability, are also possible in the agent's assessment.

True domination, or lack thereof, is a basic requirement for the plausibility of any action sequence in our model. It is akin to a sequential notion of rationalizability for information design. Following the motivation of Examples 2 and 3 in the introduction, we now seek to provide a precise characterization of truly dominance that is bereft of a search for all possible information structures.

3 True dominance: a characterization

A deviation plan is an adapted mapping $D : A \to \Delta(A)$, where recollect that being adapted means that the marginal distribution on A^t, the set of random deviation strategy for the first t periods, depends only A^t, the first t elements of the original strategy from which the agent is deviating.

Definition 5. *A deviation plan D is an* improvement over $\mathbf{a} \in A$ *if*

1. $u\left(D\left(\mathbf{b}\right), \omega\right) \geqslant u\left(\mathbf{b}, \omega\right)\ \forall\, \mathbf{b} \in A$ *and* $\omega \in \Omega$*;*
2. $u\left(D\left(\mathbf{a}\right), \omega\right) > u\left(\mathbf{a}, \omega\right)\ \forall\ \omega \in \Omega$.

Now, we can state the main result.

Theorem 1. *An action sequence* $\mathbf{a} \in A$ *is truly dominated if and only if there exists a deviation plan* D *that is an improvement over* $\mathbf{a} \in A$.

The "if" part is relatively easy, and follows mechanically from definitions. The "only if" part is not obvious. It invokes the *minmax theorem*; minimum over nature's action in choosing an adversarial prior and information structure that puts a positive weight on the truly dominated action sequence, and maximum over an alternate strategy for the agent.

The application of minmax theorem requires a two key subtle arguments. First, the set of all priors and information structures that put a positive mass on one action sequence is not compact. So we have to consider the set which puts a small minimum weight ε on the truly dominated action and then take this ε to zero. Second, we also exploit the following *revelation principle*. It is without loss of generality to assume that nature chooses an information structure which recommends an action (viz. $\alpha \circ \pi$), and then simply follows that recommendation. The agent's strategy involves taking that recommended action and choosing another action-a deviation plan.

4 Towards a theory of apparent mistakes

The fixed belief model is fairly standard in economics. For example, in a survey of *Psychology and Economics*, DellaVigna [2009] presents a version of the model to motivate departures from

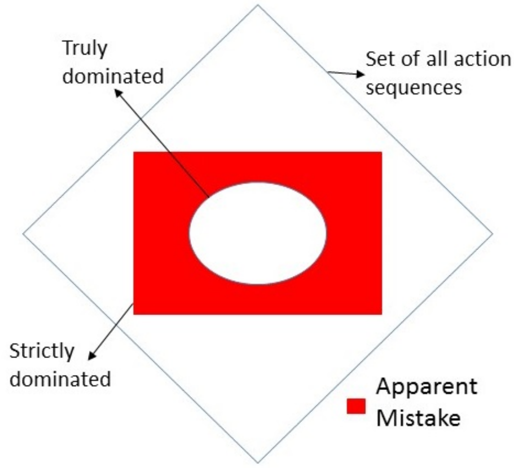

Figure 2: Strictly dominated action sequences that are not truly dominated can be construed to be mistakes

it as behavioral biases. We propose the gradual learning model as a way to understand the set of all actions taken over time by a Bayeisan agent that can be rationalized by an outside observer or analyst. As Figure 2 exposits, the set of actions that are strictly dominated but not truly dominated may be construed to be "irrational" or "suboptimal" under a fixed belief model; we call them *apparent mistakes*. This provides a novel theory that can potentially explain the choice of seemingly dominated actions in situations where time plays an important role.

With a slight strengthening of assumption on the prior, one can say a bit more. For any fixed prior, the analyst can assign a maximum probability with which a strictly dominated action sequence can be picked by a Bayesian agent; this result is stated as a corollary.

Corollary 1. *Fix a strictly dominated action sequence* $\mathbf{a}$ *that is not truly dominated, and a prior* p*. There exists* $0 < \bar{\mu} < 1$*, function of payoffs and prior, such that for any sequential information structure* (π, S) *the optimal strategy of agent,* α^**, must satisfy* $\sum_{\omega \in \Omega} \alpha^* \circ \pi(\mathbf{a}|\omega) p(\omega) \leq \bar{\mu}$.

Suppose the analyst is looking at a large population of agents making the same dynamic decision. Suppose further that either (i) each agents receives an iid prior such that the average of these mimics the actual data generating process, or (ii) each agent has a state assigned by nature unknown to her, and the aggregate distribution of these is the prior of all agents.[3] Under either of those assumptions, our analyst can also try to make sense of cross-sectional data. The choice of a strictly dominated can be rationalized only up to a fraction $\bar{\mu}$ of the population. No matter what signals each agent may observe, there exists an upperbound on the fraction of people that may choose this sequence of actions.

5 Appendix

Proof of Theorem 1. The "if" direction is easy: Suppose that D is an improvement over $\mathbf{a} \in A$. Let $p \in \Delta(\Omega)$ be a prior distribution, $\pi : \Omega \to \Delta(S)$ be an information structure, and $\alpha : S \to \Delta(A)$

[3]Both the assumptions are akin to requiring rational expectations on the prior.

be a strategy such that $\sum_{\omega\in\Omega} \alpha\circ\pi\left(\mathbf{a}|\omega\right)p\left(\omega\right) > 0$. Since D is an adapted mapping, the composition $\beta = D\circ\alpha$ is also adapted and is therefore a valid strategy (see de Oliveira [2016] for a formal statement). Now we can show that the decision maker strictly benefits from switching from the strategy α to the strategy β:

$$
\begin{aligned}
U\left(\alpha,\pi,p\right) &= \sum_{\mathbf{b}\in A\setminus\{\mathbf{a}\}}\sum_{\omega\in\Omega} u\left(\mathbf{b},\omega\right)\alpha\circ\pi\left(\mathbf{b}|\omega\right)p\left(\omega\right) \\
&\quad + \sum_{\omega\in\Omega} u\left(\mathbf{a},\omega\right)\alpha\circ\pi\left(\mathbf{a}|\omega\right)p\left(\omega\right) \\
&< \sum_{\mathbf{b}\in A\setminus\{\mathbf{a}\}}\sum_{\omega\in\Omega} u\left(D\left(\mathbf{b}\right),\omega\right)\alpha\circ\pi\left(\mathbf{b}|\omega\right)p\left(\omega\right) \\
&\quad + \sum_{\omega\in\Omega} u\left(D\left(\mathbf{a}\right),\omega\right)\alpha\circ\pi\left(\mathbf{a}|\omega\right)p\left(\omega\right) \\
&= \sum_{\mathbf{b}\in A}\sum_{\omega\in\Omega} u\left(D\left(\mathbf{b}\right),\omega\right)\alpha\circ\pi\left(\mathbf{b}|\omega\right)p\left(\omega\right) \\
&= \sum_{\mathbf{b}\in A}\sum_{\omega\in\Omega} u\left(\mathbf{b},\omega\right)D\circ\alpha\circ\pi\left(\mathbf{b}|\omega\right)p\left(\omega\right) \\
&= \sum_{\mathbf{b}\in A}\sum_{\omega\in\Omega} u\left(\mathbf{b},\omega\right)\beta\circ\pi\left(\mathbf{b}|\omega\right)p\left(\omega\right) \\
&= U\left(\beta,\pi,p\right)
\end{aligned}
$$

Now we prove the "only if" direction. Let $\epsilon > 0$ be arbitrarily small. If $\mathbf{a}\in A$ is truly dominated, then for any prior p, information structure π and strategy α such that $\sum \alpha\circ\pi\left(\mathbf{a}|\omega\right)p\left(\omega\right) \geqslant \epsilon$, there exists an alternative strategy β that gives a strictly higher payoff. We can write this in a min-max formulation:

$$
\min_{\substack{p,\pi,\alpha \\ \sum \alpha\circ\pi\left(\mathbf{a}|\omega\right)p\left(\omega\right)\geqslant\epsilon}} \max_{\beta} U\left(\beta,\pi,p\right) - U\left(\alpha,\pi,p\right) > 0
$$

We can interpret this formulation as a statement about a zero-sum game that an analyst plays with nature: Nature chooses (p,π,α) and the analyst chooses β. Notice that

$$
\alpha\circ\pi = \alpha'\circ\Pi' \Rightarrow U\left(\alpha,\pi,p\right) = U\left(\alpha',\pi',p\right)
$$

So we can write the payoff as $V\left(\alpha\circ\pi,p\right)$. Let $\Lambda_\Pi = \{\lambda:\Omega\to\Delta\left(A\right)|\lambda = \beta\circ\pi \text{ for some } \beta\}$. Then we can rewrite the problem as

$$
\min_{\substack{p,\pi,\alpha \\ \sum \alpha\circ\pi\left(\mathbf{a}|\omega\right)p\left(\omega\right)\geqslant\epsilon}} \max_{\lambda\in\Lambda_\pi} V\left(\lambda,p\right) - V\left(\alpha\circ\pi,p\right) > 0.
$$

Therefore, among the pairs (π',α') such that $\alpha'\circ\pi' = \alpha\circ\pi$, nature would like to choose the one that minimizes the size of the set Λ_π. This can be done precisely by choosing $\pi' = \alpha \circ \pi$ and $\alpha' = Id_A$. This result can be stated as the following *revelation principle*: It is without loss of generality to assume that nature chooses an information structure which recommends an action (viz. $\alpha \circ \pi$), and then simply follows that recommendation. The agent's strategy involves taking that recommended action and choosing another action-a deviation plan. We can therefore rewrite the problem as

$$\min_{\substack{p,\lambda \\ \sum \lambda(\mathbf{a}|\omega) p(\omega) \geqslant \epsilon}} \max_D V(D \circ \lambda, p) - V(\lambda, p) > 0.$$

We can simplify the problem a bit further by noting that choosing p and λ is equivalent to choosing some joint distribution on the set $A \times \Omega$. We can then write the problem as

$$\min_{\substack{\gamma \in \Delta(A \times \Omega) \\ \gamma(\mathbf{a}) \geqslant \epsilon}} \max_D \mathbb{E}_\gamma \left[u(D(a), \omega) - u(a, \omega) \right] > 0.$$

Now note that the set $\Gamma = \{\gamma \in \Delta(A \times \Omega) \,|\gamma(\mathbf{a}) \geqslant \epsilon\}$ and the set of all adapted mappings $D : A \to \Delta(A)$ are convex and compact, and the payoff function in the game above is bilinear in γ and D. Hence we can use the min-max theorem to conclude that

$$\max_D \min_{\substack{\gamma \in \Delta(A \times \Omega) \\ \gamma(\mathbf{a}) \geqslant \epsilon}} \mathbb{E}_\gamma \left[u(D(a), \omega) - u(a, \omega) \right] > 0.$$

Let $D^*(\epsilon)$ denote the set of solutions to the problem above. $D^*(\epsilon)$ is compact and nonempty. Moreover, we can show that it is increasing in ϵ: Let $\epsilon' \geqslant \epsilon$. If $D^* \in D^*(\epsilon)$ then it must be the case that

$$\mathbb{E}_\gamma \left[u(D^*(a), \omega) - u(a, \omega) \right] \geqslant \mathbb{E}_\gamma \left[u(D(a), \omega) - u(a, \omega) \right]$$

for every $\gamma \in \Delta(A \times \Omega)$ such that $\gamma(\mathbf{a}) \geqslant \epsilon$. Clearly, the same inequality must then hold for those γ that satisfy $\gamma(\mathbf{a}) \geqslant \epsilon'$, which means that $D^* \in D^*(\epsilon)$. Hence the intersection $\bigcap_{\epsilon>0} D^*(\epsilon)$ is not empty, since it is an intersection of nested compact sets.

Let D^* be an element of that intersection. Then, $\mathbb{E}_\gamma \left[u(D^*(a), \omega) - u(a, \omega) \right] > 0$ for any $\gamma \in \Delta(A \times \Omega)$ such that $\gamma(\mathbf{a}) > 0$. Letting $\gamma = \delta_{(\mathbf{a},\omega)}$, we see that

$$u(D^*(\mathbf{a}), \omega) > u(\mathbf{a}, \omega) \qquad \forall \omega \in \Omega.$$

Letting $\gamma = \epsilon \delta_{(\mathbf{a},\omega)} + (1-\epsilon)\, \delta_{(\mathbf{b},\omega)}$, we see that

$$\epsilon u(D^*(\mathbf{a}), \omega) + (1-\epsilon) u(D^*(\mathbf{b}), \omega) > \epsilon u(\mathbf{a}, \omega) + (1-\epsilon)\, u(\mathbf{b}, \omega) \qquad \forall \epsilon > 0, \forall \omega \in \Omega \text{ and } \forall \mathbf{b} \in A.$$

Letting $\epsilon \to 0$, we get that $u(D^*(\mathbf{b}), \omega) \geqslant u(\mathbf{b}, \omega)$ for every $\mathbf{b} \in A$ and $\omega \in \Omega$. Therefore we have that D^* is an improvement over $\mathbf{a} \in A$. □

References

L. M. Ausubel. The failure of competition in the credit card market. *American Economic Review*, 81(1):50–81, 1991.

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. Yale University and Princeton University, 2017.

H. de Oliveira. Blackwell's informativeness theorem using category theory. Pennsylvania State University, 2016.

S. DellaVigna. Psychology and economics: evidence from the field. *Jounral of Economic Perspectives*, 47(2):315–372, 2009.

S. DellaVigna and U. Malmendier. Paying not to go to the gym. *American Economic Review*, 96 (3):694–719, 2006.

J. Ely. Beeps. *American Economic Review*, 107(1):31–53, 2017.

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

M. Osborne and A. Rubinstein. *A course in game theory*. The MIT Press, 1994.