Ambiguity, Learning, and Raiffa's critique

Filippo Massari

School of Banking and Finance, UNSW

February 15, 2018

Abstract

We present evidence of non-trivial interactions between learning and ambiguity. We find that ambiguous averse decision functional do not produce robust out-of-sample predictions when matched with learning models. There are situations in which the illusion of learning induces an ambiguous averse DM to optimally choose a sequence of ambiguous acts over a sequence of risky acts which would deliver a higher average utility.

1 Introduction

In every period a DM can chose to bet on the color of a ball extracted from either a risky or an ambiguous urn (Ellsberg, 1961). The DM faces a version of the standard dilemma of exploration vs exploitation: while the risky urn is fair by construction, the ambiguous process might be favorable for bets on one of the colors. Luckily, our DM heard about Raiffa (1961)'s critique and knows that he can use a coin to "explore" the ambiguous alternative while hedging away its ambiguity.[1]

[1] Raiffa (1961): *"Suppose you withdraw a ball from the urn with unknown composition but do not look at its color. Now toss a fair (unbiased) coin and call "red" if heads, "black" if tails. The "objective" probability of getting a match is now .5, and therefore it is just as desirable to participate in the second game (ambiguous urn) as in the first (risky urn).*"

Is it "safe" for the DM to take advantage of this opportunity?

We give negative answer to this question. Consider the following strategy, *Raiffa strategy* (RS). The DM bets on the ambiguous urn in every period. In those periods in which he faces ambiguity, he uses a coin to objectivize it — in this way he samples from the ambiguous process while facing risky payoffs. The DM does not use the coin in those periods (if any) in which empirical evidence has persuaded the DM that betting on a certain color in the ambiguous urn ensures higher unambiguous expectation than betting on the risky urn.

Our analysis of the RS shows that learning can work as a double-edge sword. RS is potentially profitable because it allows the DM to learn about profitable ambiguous processes. However, there are also situations in which RS delivers strictly lower average payoffs than betting on the risky urns. Depending on the path of realizations, we identify three possible cases.

1) The RS prescribes to bet on the same color a unitary fraction of periods. In this case, the RS delivers a (weakly) higher average utility than betting on the risky urns.[2]

2) The RS prescribes to randomizes a positive but non-unitary fraction of periods. In this case, the RS delivers a strictly lower average utility than betting on the risky urns.

3) The RS prescribes to randomizes a unitary fraction of periods. In this case, the RS delivers the same average utility as betting on the risky urns.

Case 1 describes situations in which the DM eventually learns unambiguously that one color (e.g., white) has higher average payoff than the other in the ambiguous process. By always betting on white the DM achieve weakly higher average payoff than randomizing. In case 2, the DM has lower average utility following RS than betting on the risky urns because whenever he chose a color — because he believes to have learned something useful about the ambiguous data generating process—, he is later proven wrong by the data. This result is a generalization of the well-known

[2]This is the only relevant case if balls are extracted with replacement from the same ambiguous urn in every period.

fragility of using in sample estimates to make out-sample predictions when the model is misspecified. It shows that even if ambiguous decision functional are more conservative than standard statistical procedures, they do not deliver predictions that are robust regarding out-of-sample prediction accuracy. Case 3 describes the situation in which the empirical evidence never supports a color strongly enough to persuade the DM to bet on it. The DM randomizes in every period, and his average payoff is the same as betting on the risky urn.

Our result exhorts caution in these situations in which a decision criterion prescribes switching discontinuously between optimal actions. For example, it formalizes the intuition that using different stock evaluation criteria (e.g., different valuation from fundamental analysis, opinions of a pool of experts, or empirical methods) and deciding to buy and hold a stock only if all criteria indicate that the stock is undervalued — as routinely recommended in portfolio evaluation textbooks — is a dangerous practice because it generates sub-optimal out of sample returns in all those cases in which the "best" stock changes over time — e.g., because stock returns follow a mean-reverting process.

2 Learning

We model inter-temporal ambiguity as the scenario in which a Bayesian learner holds more than one prior distribution over a set of parameters (Marinacci, 2002; Marinacci and Massari, 2017).[3] We consider a family of iid[4] models $\mathcal{M} = \{P_\theta : \theta \in \Theta\}$ with a parameter set $\Theta \subset \mathbb{R}^n$, defined on a σ-algebra Σ^∞ of subsets of X^∞ with representative element $x^\infty = x_1, x_2, ...$; where $X^\infty := \times^\infty X$ is the infinite cartesian product of a finite observation space X with representative element x and σ-algebra Σ.[5] With a slight

[3]We use the multiple prior setting for illustration. The general result — that objective randomization renders the ambiguous setting payoff equivalent to the risky setting if and only if the DM relies on randomization a unitary fraction of periods — does apply to other learning environments (e.g., Epstein and Schneider, 2007; Epstein and Seo, 2015).

[4]In Section 7.3 we generalize the result to the non-iid setting.

[5]In this paper, we focus on extractions from ambiguous urns. However, this setting can accommodate most prediction tasks with minor changes which do not affect our results. For example, x could be a vector of stock market returns, $\mathcal{M}$ a set of regression models with parameters to be estimated and $\mathcal{C}$ a (meta)prior over the set of regression.

abuse of notation, we use $P_\theta(x^t)$ to denote the probability that model P_θ attaches to the cylinder with base x^t, and the likelihood that model P_θ attaches to the partial sequence $(x_1, ..., x_t)$. The prior information about the parameters is summarized by prior distributions $\mu : 2^\Theta \to [0,1]$. The set of prior distributions is $\mathcal{C}$. For any prior distribution $\mu \in \mathcal{C}$ the joint distribution of the parameters and the observations is $P_\mu : 2^\Theta \to [0,1]$. By definition, for all $A \subseteq \Theta$ we have that:

$$P_\mu(A \times x^t) := \int_A P_\theta(x^t) d\mu.$$

We denote by $\mu(.|x^t) : 2^\Theta \to [0,1]$ the usual posterior given the observations x^t, while $P_\mu(.|x^t) : \Sigma \to [0,1]$ is the one step ahead predictive distribution of x_{t+1}, given observations x^t. By definition, for all $A \subseteq \Theta$:

$$P_\mu(A \times x_{t+1}|x^t) := \int_A P_\theta(x_{t+1}) d\mu(.|x^t) := \int_A P_\theta(x_{t+1}) \frac{P_\theta(x^t) d\mu}{\int_\Theta P_\theta(x^t) d\mu}.$$

3 Preferences

Let C be the space of consequences on which the DM has a bounded utility function $u : C \to \mathbb{R}$. An act $f : X \to C$ is a $\Sigma-$measurable map that associates a consequence to each observation in X. $\mathcal{F}$ is the space of available acts for the DM. We are considering one step ahead acts. The decision criterion adopted by the DM depends on the quality of his prior information. In evaluating an act in this scenario, the DM has to use a set criterion. In the first period, the DM's set criterion is for every act f, given by:

$$\left\{ \int_\Omega u(f(x)) dP_\mu(.|\emptyset) : \mu \in \mathcal{C} \right\}.$$

Subsequently, as the DM incorporates past realization, x^t, to each prior in $\mathcal{C}$ using Bayes' rule, his choice criterion becomes:

$$\left\{ \int_\Omega u(f(x)) dP_\mu(.|x^t) : \mu \in \mathcal{C} \right\}.$$

For illustrative purposes,[6] we assume that the DM preferences are described by maxmin criterion of Gilboa and Schmeidler (1989).

$$\forall t, f \succ g \Leftrightarrow \min_{\mu \in \mathcal{C}} E_{\mu,t} u(f(x_t)) \geq \min_{\mu \in \mathcal{C}} E_{\mu,t} u(g(x_t)),$$

where $E_{\mu,t}$ is the expectation according to $P_\mu(\cdot|x^{t-1})$.

4 Raiffa strategy and act-relevance

To incorporate Raiffa recommendation in our decision problem, we enrich the DM action space with an unambiguous act f^{r^*}. Act f^{r^*} represents the lottery obtained by choosing an act at random after the state has realized but before it has been observed. As highlighted by Raiffa, f^{r^*} effectively transforms an ambiguous setting into a risky one. In the standard urn setting, act f^{r^*} consists on tossing a coin to decide which color to bet on. We say that

Definition 1. *A DM follows* Raiffa Strategy *(RS) if in every period t after history x^{t-1}, he choses act*

$$f_t^*(x^{t-1}) := \underset{i \in \mathcal{F} \cup f^{r^*}}{\operatorname{argmax}} \left\{ \min_{\mu \in \mathcal{C}} E_{\mu,t-1} u(f_i(x_t)), E_{r^*} u(f^{r^*}(x_t)) \right\}$$

where $E_{\mu,t}$ is the expectation according to $P_\mu(\cdot|x^{t-1})$ and E_{r^*} is the expectation according to the objective randomization device.

This strategy describes the optimal choice of a DM with maxmin preferences, multiple prior learning and whose action set also includes the possibility to randomize: f^{r^*}. In every period, the RS prescribes to randomize if and only if the expected utility of the risky bet is higher than the expected utility according to the least favourable posterior obtained from priors in $\mathcal{C}$.

[6] Our results generalizes streightforwardly to other decision functionals — including SEU (Savage, 1954), KMM (Klibanoff et al., 2005); VP (Maccheroni et al., 2006) —, and to decision functionals describing incomplete preferences with inertia (Aumann, 1962; Walley, 1991; Bewley, 2002).

The performance of RS critically depends on whether or not a unique act is recommended a positive fraction of periods.We say that ambiguity is *act-relevant* on those sequences in which DM that follows RS change its action a positive fraction of periods. While ambiguity is *act-irrelevant* if a DM that follows RS chooses the same action in most periods.

Definition 2. *Ambiguity is* act-relevant *on path x^∞ if,* $\lim_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} I_{\{f^*_\tau \neq f^*_{\tau-1}\}} > 0$; *Ambiguity is* act-irrelevant *on path x^∞ if,* $\lim_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} I_{\{f^*_\tau \neq f^*_{\tau-1}\}} = 0$.[7]

There are two ways for ambiguity to be act-irrelevant on a path. Either learning does not occur, and RS prescribes to randomize in most periods. Or the DM eventually learns that one of the non-random acts delivers higher utility than the risky act and RS prescribes it for all large t.

In our multiple-prior setting, ambiguity can be act-relevant only if the learning problem is misspecified. In other words, if none of the models in the support of the DM are correct. While this occurrence is a tight requirement from the theoretical perspective — it requires having at least two models in $\mathcal{P}$ with the same K-L divergence. We stress that, in practical prediction problems, having more than a valid candidate model is the norm rather than the exception because competing models are not selected at random (Marinacci and Massari, 2017). This situation is known to arise often in the finance literature (Avramov, 2002; Cremers, 2002), and it has stimulated the growing interests in predictions based on model averages with weights that do not evolve according to Bayes' rule (Timmermann, 2006).

5 Main result

In this Section, we discuss the average utility that a DM obtains by betting on a sequence of ambiguous urns following RS and compare it with the average utility obtained by betting on the risky urns.

[7]We assume that these limits exists.

Theorem 1. *Let* $\bar{u}(f^*(x^\infty)) = \lim_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} u(f^*_\tau(x^{\tau-1}))$*, and* $\bar{u}(f^r(x^\infty)) = \lim_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} u(f^r)$*, be the average utilities that a DM obtains on path* x^∞ *using RS and investing in the risky urns, respectively. The following holds:*

i) $\bar{u}(f^*) \geq \bar{u}(f^r)$ P*-a.s., in those sequences in which ambiguity is act-irrelevant and* $f^*_t \neq f^{r^*}$ *in most periods;*

ii) $\bar{u}(f^*) = \bar{u}(f^r)$ $P^* \times P$*-a.s., in those sequences in which ambiguity is act-irrelevant and* $f^*_t = f^{r^*}$ *in most periods;*

iii) $\bar{u}(f^*) < \bar{u}(f^r)$ $P^* \times P$*-a.s., in those sequences in which the long-run average exists and ambiguity is act-relevant.*

Where P *and* P^* *are the probability generated by the risky urn and the randomization device, respectively.*

Point *i*) shows that if an act becomes unambiguously preferred to randomizing a unitary fraction of periods, then the DM's average utility using RS is at least as high as what he would obtain by betting on the sequence of risky urns. The inequality is strict in those sequences in which one of the colors in the ambiguous urn has a higher frequency than the other; while equality holds in those sequences in which the frequency of a color approaches .5 monotonically for all large t.

Point *ii*) tells us that in those sequence in which ambiguity is act-relevant the DM has lower average utility adopting RS in the sequence of ambiguous urns rather than betting on the risky urns. Because the optimal act changes a positive fraction of periods, there is a positive fraction of periods in which a deterministic mismatch between the act chosen by the DM and the realized outcome occurs. In these periods the DM regret having chosen the color deterministically rather than using the coin. The RS delivers a lower average utility than the risky urns when ambiguity is act relevant because his average utility is not higher than that he would achieve in the risky urns in those periods in which he randomizes, and in those periods in which he does not change act, while it is strictly lower when there is a mismatch.

Point *iii*) follows from the Strong Law of Large Numbers. It tells us that if the DM chooses to randomize a unitary fraction of periods, then his average payoff equal to what he would obtain by betting on the risky urns.

Next, we present an example to illustrate the different cases of Theorem 1.

Example 1 Consider a DM facing a series of ambiguous urns with balls of two colors: White (x_t^w) and Black (x_t^b).

Acts: $\forall t$, the DM can choose between three acts $\begin{cases} f_t^b: \text{to bet on a black ball;} \\ f_t^b: \text{to bet on a white ball;} \\ f_t^{r^*}: \text{to use a fair coin to chose between } f^b \text{ and } f^w. \end{cases}$

Consequence: $\forall t$, the DM Payoffs are : $\begin{array}{ccc} & x_t^b & x_t^w \\ f_t^b & \$100 & 0 \\ f_t^w & 0 & \$100 \end{array}$.

Beliefs: The DM believes realisations are iid from two possible models: $\Theta = \{\theta_1; .\theta_2\}$.

Where, $\forall t$, $\theta_1 = Pr(x_t = b) = .2$ and $\theta_1 = Pr(x_t = b) = .8$.

The DM has two priors on Θ: $\mathcal{C} := \left\{\mu^1, \mu^2\right\}$.

Where $\mu^1(\theta_1) = .1 = 1 - \mu^2(\theta_2)$.

Theorem 1 i) applies (with strict inequality) in all those sequences in which the average number of black (white) exceeds that of white (black), $x^\infty : \lim\limits_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^t I_b = \bar{t}_b > .5$.

Proof:

$\bar{t}^b > .5 \Rightarrow$ the posteriors from both prior converge to $\theta = .8$: $\begin{cases} P(x_t^b|\mu_1) \to .8 \\ P(x_t^b|\mu_2) \to .8 \end{cases}$

$\Rightarrow$ the DM chooses black, a unitary fraction of periods: for most t, $f_t^* = f^b$

$\Rightarrow \bar{u}(f^*) = \bar{t}_b u(\$100) + (1 - \bar{t}_b)u(0) > .5u(\$100) + .5u(0) =^{P\text{a.s.}} \bar{u}(f^r)$.

Theorem 1 ii) is relevant for the sequence $x^\infty := \{B, B, W, B, W....\}$.

Proof:

$$x^\infty = \{B,B,W,B,W....\} \Rightarrow \begin{cases} \min\{P(x_t^b|\mu_1), P(x_t^b|\mu_2)\} > .5 & \forall \text{ odd } t \geq 3 \text{ ,} \\ \begin{cases} \min\{P(x_t^w|\mu_1), P(x_t^w|\mu_2)\} < .5 \\ \max\{P(x_t^w|\mu_1), P(x_t^w|\mu_2)\} > .5 \end{cases} & \forall t \text{ even.} \end{cases}$$

$$\Rightarrow \text{The DM chooses:} \begin{cases} f_t^* = f^b & \forall \text{ odd } t \geq 3 \text{ ,} \\ f_t^* = f^{r^*} & \forall t \text{ even.} \end{cases}$$

$$\Rightarrow \bar{u}(f^*(x^\infty)) = O(\frac{1}{t}) + \lim_{t\to\infty} \frac{1}{t}\sum_{\tau=3}^{t} u(f_t^*)$$

$$= \lim_{t\to\infty} \frac{1}{t}\sum_{\tau=3}^{t} u(f_{u,t}(x))\bigg|_{\tau\ odd} + \lim_{t\to\infty} \frac{1}{t}\sum_{\tau=3}^{t} u(f_{u,t}(x))\bigg|_{\tau\ even}$$

$${}^{a} =^{P^*\text{a.s.}} .5u(0) + .5(.5u(0) + .5u(100))$$

$$< .5u(0) + .5u(100)$$

$$=^{P\text{a.s.}} \bar{u}(f^r)$$

a) The $.5u(0)$ term reflects the fact that the DM choses f^b incorrectly in every $\tau \geq 3$ odd; while the $.5(.5u(0)+.5u(100))$ term follows from the SLLN because in every even period with $\tau \geq 4$ the DM has probability .5 to guess correctly the color since he is randomizing.

Theorem 1 iii) is relevant for the sequence $x^\infty := \{W,B,W,B,....\}$.

Proof:

$$x^\infty := \{W,B,W,B,....\} \Rightarrow \begin{cases} \min\{P(x_t^b|\mu_1), P(x_t^b|\mu_2)\} < .5 \\ \max\{P(x_t^b|\mu_1), P(x_t^b|\mu_2)\} > .5 \end{cases} , \forall t$$

$\Rightarrow$ the DM randomizes a unitary fraction of periods: $\forall t, f_t^* = f_t^{r^*}$.

$\Rightarrow \bar{u}(f^*) =^{P^*\text{-a.s.}} .5u(\$100) + .5u(0) =^{P\text{-a.s.}} \bar{u}(f^r)$.

6 Conclusion

Ambiguous averse decision functional, when matched with learning models do not produce robust out-of-sample predictions. There are situations in which the illusion of learning induces an ambiguous averse DM to optimally choose a sequence of ambiguous acts over a sequence of risky acts which would deliver a higher average utility.

7 Discussion

7.1 Act-relevance vs long-run ambiguity

Act-relevance is closely related to the notion of long-run ambiguity proposed and analyzed by Marinacci and Massari (2017) — A DM is said to suffer long run ambiguity if the posteriors obtained from all priors in $\mathcal{C}$ do not converge to a Dirac measure on the same model in $\mathcal{M}$. Here we show that act-relevance and long-run ambiguity, although closely related, are not equivalent notions. There are cases in which ambiguity fades away, but a unique action does not appear, and cases in which ambiguity does fade away and a unique best action appears.

Let start with some clarifying definitions.

Definition 3. *Ambiguity fades away at path $x^\infty \in X^\infty$ if, for every act f,*

$$\lim_{t\to\infty}\left[\sup_{\mu\in\mathcal{C}}\int_X u(f(x))dP_\mu(.|x^t) - \inf_{\mu\in\mathcal{C}}\int_X u(f(x))dP_\mu(.|x^t)\right] = 0;$$

where, $\forall t > 0, x^t$ indicates the first t realizations of path x^∞.

Marinacci and Massari (2017) show that ambiguity fades away if and only if the data clearly support a best model among those believed possible by the DM. Where "best model" formally means:

Definition 4. *Given a path $x^\infty \in X^\infty$ and a family of models $\mathcal{M} = \{P_\theta : \theta \in \Theta\}$, with $|\Theta| < \infty$. We say that $\hat{\theta} := \hat{\theta}(x^\infty, \Theta)$ is the* strong maximum likelihood parameter *if $\hat{\theta} \in \Theta$ and*

$$\forall\theta \in \Theta, \theta \neq \hat{\theta} \Rightarrow \lim_{t\to\infty}\frac{P_\theta(x^t)}{P_{\hat{\theta}}(x^t)} = 0;$$

where, $\forall t > 0, x^t$ indicates the first t realizations of path x^∞.

Their main theorem demonstrates that the existence of a *strong maximum likelihood parameter* is a necessary and sufficient condition for all posteriors to concentrate on the same parameter and for ambiguity to fade away.

Theorem. *(Marinacci and Massari, 2017): let $\mathcal{M} = \{P_\theta : \theta \in \Theta\}$ be a (finite) family of models and $\mathcal{C}$ a compact set of non-degenerate prior distributions on Θ. Then, ambiguity fades away at path x^∞ if and only if $\hat{\theta}(x^\infty, \Theta)$ exists.*[8]

The fading away and the act-relevance definitions are closely related, but not equivalent. While the former is a statement about the next period expected utility, the latter is a statement about the actual choice made by the DM. The next theorem shows that these definitions are equivalent if and only if the utility function is continuous. Otherwise the relative strength of the two definitions depends on the cardinality of Θ.

Theorem 2. *Let $|\Theta| < \infty$, then*

a) ambiguity vanishes on $x^\infty \Rightarrow$ ambiguity is act-irrelevant on x^∞;

b) ambiguity is act irrelevant on $x^\infty \not\Rightarrow$ ambiguity vanishes on x^∞;

c) ambiguity is act irrelevant on x^∞ and u is continuous, non-constant $\Rightarrow$ ambiguity vanishes on x^∞;

Let $|\Theta| = \mathbb{R}$, then

a') ambiguity is act-irrelevant on $x^\infty \Rightarrow$ ambiguity vanishes on x^∞;

b') ambiguity vanishes on $x^\infty \not\Rightarrow$ ambiguity is act-irrelevant on x^∞;

c') ambiguity vanishes on x^∞ and u is continuous, non-constant $\Rightarrow$ ambiguity is act-irrelevant on x^∞;

Theorem 2 confirms that act relevance and the vanishing condition are close notions. However, its points b) and b') highlight that act-relevance and long run ambiguity are not equivalent when $u(\cdot)$ is not continuous, and it shows how the relative strength of the two conditions depends on the cardinality of Θ. If $|\Theta| < \infty$ and $u(.)$ is discontinuous, the ambiguity vanishing condition is stronger than act-irrelevant. There are cases in which ambiguity does not vanish — the posteriors do not concentrate on the same model—, and yet a unique best action is chosen for all large t. Example 2 illustrates a case in which ambiguity is act-irrelevant even if it does not fade away.

[8] For the only if part of the result we implicitly assumed that for all acts f and for all $x, x', x \neq x' \Rightarrow u(f(x)) \neq u(f(x'))$.

Example 2: in the same setting of Example 1, consider the decisions made by a DM who follows RS on the deterministic sequence $\mathrm{x}^\infty := (B, B, B, W, B, W, B, ...)$. Clearly, on x^∞ ambiguity does not fade away because the posteriors do not converge to a Dirac in any model. However, ambiguity is not act relevant because $\forall t \geq 2, f_t^* = f^b$.

On the other hand, if $|\Theta| = |\mathbb{R}|$ and if $u(.)$ is discontinuous, it is the act-relevance condition that it is harder to satisfy. When u is discontinuous, asymptotic relevance is stronger than the vanishing condition because an infinitesimal disagreement on the posteriors suffices to determine a different optimal action. Example 3 illustrates a case in which ambiguity fades away even if it is not act-irrelevant.

Example 3: in the same setting of Example 1, suppose that the DM subjectively believes that realizations are iid and that all compositions of the urn are possible. Formally, he believes that $\forall t, \theta_t = \theta \in \Theta = (0, 1)$ where θ the number of black balls in the urn. The prior information of the DM on the composition of the urn is represented by the family of Beta priors with parameters in two strictly positive, finite intervals: $\mathcal{C} = \{Beta(\alpha, \beta),\ \alpha \in (0, \infty), \beta \in (0, \infty)\}$ on Θ, which We need to show that there are sequences in which ambiguity fades away, and yet it remains act-relevant.

- Ambiguity fades away: Marinacci and Massari (2017) (online appendix) show that ambiguity fades away in all sequences.
- There are sequences in which ambiguity remains act-relevant: it is easy to verify that ambiguity is act-relevant, among others, in all sequences x^∞ such that

$$\begin{cases} \lim\limits_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^t I_{\{\bar{b}(x^\tau) > .5\}} > 0 \\ \lim\limits_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^t I_{\{\bar{b}(x^\tau) < .5\}} > 0 \end{cases},$$

  where $\bar{b}(x^t) := \min_{\alpha,\beta} \frac{\alpha+\tau_r}{\tau+\alpha+\beta}$ is the expected number of black balls according to the most conservative posterior obtained from the priors in $\mathcal{C}$. For example, it can be verified that if $\alpha, \beta \in [1, 2]$ any mean reverting process with frequency of black equals to .5 satisfies this condition.

7.2 Raiffa Strategy Vs Fictitious Player Dynamic

In this section, we draw a parallel between the performance of the RS and that of the Fictitious Player Dynamic, FPD, (Brown, 1951; Fudenberg and Levine, 1998). If our $\mathcal{C}$ were a singleton containing only the uniform distribution, the optimal decisions according to the RS would coincide with the best responses according to the FPD. In every period, the RS (FPD) prescribes to chose the act that is the best response to the frequency of past realizations (action of the opponent). As per FPD, the sequences in which the RS underperforms against randomizing in every period — the min-max Nash equilibrium — are those in which there are infinitely many changes in best responses — which happens when the data generating process is not iid as erroneously assumed by the DM (player). When $\mathcal{C}$ is not a singleton, the main difference between the RS and the FPD is that the RS is more cautious than the FPD as it prescribes to randomize more often than the FPD does. While the FPD prescribes randomizing if the posterior calculated from the uniform prior renders the two choices equivalent; the RS prescribes to randomize in all those cases in which at least one of the posteriors calculated from the priors in $\mathcal{C}$ indicates that randomizing delivers an expected utility at least as high as not randomizing.

The similarities between the RS and the FPD support the conjecture that a strategy that guarantees good worst-case average utilities in our multiple prior setting should rely on a smooth transition between randomized and non-randomized decisions rather than switch deterministically between the two extremes (in the spirit of Fudenberg and Levine, 1998; Freund and Schapire, 1999; De Rooij et al., 2014). While adapting known algorithms to our multiple priors setting should not be excessively challenging, we do not go in that direction because it would be hard to motivate such a strategy from a decision theory perspective. While these algorithms treat prediction — beliefs in our setting — as a choice variable, which is chosen to minimize a loss function (the negative of a utility function). Decision theory axiomatically treats beliefs as primitives of the choice problem.

7.3 Non-iid setting

We derived our results assuming that the DM only considers iid models to be possible. This assumption eases the interpretation, but it is not needed. All our results apply verbatim if we generalize our the definition of act relevance as follows.

Definition 5. *Ambiguity is* act-relevant *on path x^∞ if,*

$$\forall P_\theta \in \mathcal{M} : \exists \mu \in \mathcal{C} \text{ and } P_{\theta'} \in \mathcal{M} : \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} I_{\left\{ \frac{P_\theta(\sigma^\tau)\mu(\theta)}{P_{\theta'}(\sigma^\tau)\mu(\theta')} > 1 \ \wedge \ f^*_\tau|_{P_\theta} \neq f^*_\tau|_{P_{\theta^*}} \right\}} \in (0,1);$$

Ambiguity is act-irrelevant *on path x^∞ if one of these two conditions is satisfied,*

$$\begin{cases} \exists P_\theta \in \mathcal{M} : \forall \mu \in \mathcal{C}, \forall P_{\theta'} \in \mathcal{M} \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} I_{\left\{ \frac{P_\theta(\sigma^\tau)\mu(\theta)}{P_{\theta'}(\sigma^\tau)\mu(\theta')} > 1 \ \vee \ f^*_\tau|_{P_\theta} = f^*_\tau|_{P_{\theta^*}} \right\}} = 1; \\ \forall P_\theta \in \mathcal{M} : \exists \mu \in \mathcal{C} \text{ and } P_{\theta'} \in \mathcal{M} : \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} I_{\left\{ \frac{P_\theta(\sigma^\tau)\mu(\theta)}{P_{\theta'}(\sigma^\tau)\mu(\theta')} > 1 \ \wedge \ f^*_\tau|_{P_\theta} \neq f^*_\tau|_{P_{\theta^*}} \right\}} = 0. \end{cases}$$

Lemma 1. *Definition 2 coincides with Definition 5 if all models in $\mathcal{M}$ are iid.*

Proof. It follows from the definition of f^*, noticing that in iid setting $f^*_t|_{P_\theta}$ does not depends on t. $\square$

Definition 5 makes the connection between act relevance and the performance of the RS more transparent. Each condition determines a different performance. It tells us that ambiguity is act-relevant if the DM changes his opinion about what is the best model a positive fraction of periods, with the caveat that this change of opinion must have implications on his actions.

Furthermore, Definition 5 makes the comparison between act-relevance and the vanishing condition more transparent. While vanishing requires the existence of a model whose likelihood dominates that of all others in the sequence, act irrelevance only requires the likelihood ratio between on model and the others to be greater than 1.

A Appendix

Proof of Theorem 1

Proof. We prove our results under the iid assumption. The proof for the non-iid case follows the same logic, but adopting Definition 5 for act-relevance.

The average utility of the DM on an arbitrary sequence x^∞ can be decomposed as follow:

$$\begin{aligned}
&\lim_{t\to\infty}\frac{1}{t}\sum_{\tau=1}^{t} u(f^*(x_\tau|x^{\tau-1})) \\
&=^a \lim_{t\to\infty}\frac{1}{t}\sum_{\tau=1}^{t} u(f^*(x_\tau|x^{\tau-1}))\Big|_{f^*_\tau=f^{r*}} && (=\alpha(.5u(100\$)+.5u(0\$))\ P^*\text{-a.s.} \\
&+^b \lim_{t\to\infty}\frac{1}{t}\sum_{\tau=1}^{t} u(f^*(x_\tau|x^{\tau-1}))\Big|_{f^*_\tau\neq f^*_{\tau+1}=f^{r*}} && (=\beta(u(0\$))) \\
&+^c \lim_{t\to\infty}\frac{1}{t}\sum_{\tau=1}^{t} u(f^*(x_\tau|x^{\tau-1}))\Big|_{f^*_\tau=f^*_{\tau+1}\neq f^{r*}} && (=\gamma(.5(u(100\$)+.5u(0\$)) \\
&+^d \lim_{t\to\infty}\frac{1}{t}\left|\sum_{\tau=1}^{t} (I_{x_\tau=w}-I_{x_\tau=b})\,u(100\$)\right| && (=\delta u(100\$))
\end{aligned}$$

Where α, β, γ are the fraction of periods in which (α) the DM randomizes, (β) he does't randomize and will change his action in the next period, and (γ) he doesn't randomize and will not change action in the next period, respectively; and δ is the average absolute distance between the number of white and black realizations— it is a measure of the intensity of the empirical support on one of the colors.

By construction $\alpha+\beta+\gamma+\delta=1$, the average payoffs the DM gets in each segment are indicated on the right of the equation and obtained as follows.

a) The first line is the average payoff that the DM gets in those periods in which he uses the randomisation device. In these periods his expected utility coincides with what he would get betting on the risky urns. Thus, by the Strong Law of Large Numbers, his average utility equals .5u(100\$)+.5u(0\$) almost surely.

b) The second line is the average payoff of those periods before the DM changes act.

The fact that he changes act at $\tau+1$ reveals that he would have preferred to have chosen a different act in τ, with hindsight. Because his choice is deterministic at τ, this means that he chose to bet on the incorrect color, which in turns, it means that his utility in those periods is u(0).

c) The third line is the average payoff of those periods in which the DM does not randomize and doesn't change strategy — because all priors agree that choosing to bet on a color (e.g., white) delivers higher expected payoffs than randomizing even after observing the actual realization. Fig 1 shows that the number of periods in which the DM choice is correct among these periods is exactly .5(u(100\$)+.5u(0\$)) plus the forth line.

d) The fourth line is the deviation between .5 and the average number of white (black) balls extracted.

Therefore, the Theorem is proven comparing

$$\lim_{t\to\infty}\frac{1}{t}\sum_{\tau=1}^{t}u(f^*(x_\tau|x^{\tau-1})) =^{P^*\text{-a.s.}} (\alpha+\gamma)(.5u(100\$)+.5u(0\$))+\beta u(0\$)+\delta(u(\$100))$$

$$\lim_{t\to\infty}\frac{1}{t}\sum_{\tau=1}^{t}u(f^r(x_\tau)) =^{P\text{-a.s.}} .5u(100\$)+.5u(0\$)$$

i): Act-irrelevance and $f^*\neq f^{r^*}\Rightarrow \bar{u}(f^*)\geq\bar{u}(f^r)$ P-a.s..

Proof: Act-irrelevance $\Rightarrow \alpha,\beta=0\Rightarrow \bar{u}(f^*)\geq\bar{u}(f^r)$ P-a.s.;

$f^*\neq f^{r^*}\Rightarrow \bar{u}(f^*)>\bar{u}(f^r)$ P-a.s. $\Leftrightarrow \delta>0$.

ii): Act-relevance and the existence of the long run average $\Rightarrow \bar{u}(f^*)<\bar{u}(f^r)$ $P^*\times P$-a.s..

Proof: Act-relevance and the existence of the long run average $\Rightarrow \delta=0:$

By contradiction, suppose $\lim\limits_{t\to\infty}\frac{1}{t}\left|\sum_{\tau=1}^{t}\left(I_{x_\tau=w}-I_{x_\tau=b}\right)u(100\$)\right|\neq 0$. This implies that the number of white (black) realizations exceed that of (white) black realizations by an order $O(t)$. If the limit exists, this violates the assumption that ambiguity is not asymptotically relevant. If the limit does not exist, then the existence of the long-run average assumption is violated (standard result in non-standard analysis).

iii): Act-irrelevance and $f^* = f^{r^*} \Rightarrow \bar{u}(f^*) \geq \bar{u}(f^r)$ $P^* \times P$-a.s..

Proof: Act-irrelevance $\Rightarrow \beta = 0$, $f^* = f^{r^*} \Rightarrow \gamma = \delta = 0$.

$\square$

Proof of Theorem 2

Proof. $|\Theta| < \infty$

a): ambiguity vanishes $\Leftrightarrow$ all posteriors converge to a Dirac on a unique model ($|\Theta| < \infty$ with orthogonal models $\Leftrightarrow$ models are not contiguous) $\Rightarrow f_t^*$ is unique for all t large.

b): Ambiguity is act-irrelevant on $x^\infty \not\Rightarrow$ ambiguity vanishes on x^∞: see example 2 for a case in which ambiguity does not vanishes but is act-irrelevant.

c): Ambiguity is act-irrelevant on x^∞ and u is continuous, non-constant $\Rightarrow$ ambiguity vanishes on x^∞:

f_t^* is constant for all large t and u continuous not constant

$\Rightarrow$ the predictive probabilities are constant for all large t

$\Leftrightarrow$ all posteriors are degenerate on the same model for all large t

$\Leftrightarrow$ ambiguity vanishes on x^∞.

$|\Theta| = |\mathbb{R}|$

a') The statement is trivially true because ambiguity vanishes on every sequence — under mild regularity condition on the set of priors (Marinacci and Massari, 2017, online appendix).

*b')*Ambiguity vanishes on $x^\infty \not\Rightarrow$ ambiguity is act-irrelevant on x^∞: see example 3 for a case in which ambiguity is not Act-irrelevant but vanishes.

c') Ambiguity vanishes on x^∞ and u is continuous, non-constant $\Rightarrow$ ambiguity is act-irrelevant on x^∞. By contradiction, suppose ambiguity is not act-irrelevant. Thus,

$\exists \epsilon > 0 : \|f_t^* - f_{t-1}^*\| > \epsilon$ infinitely often

$$\Leftrightarrow \left\| \underset{i \in \mathcal{F} \cup f^{r^*}}{\operatorname{argmax}} \left\{ \min_{\mu \in \mathcal{C}} E_{\mu,t-1} u(f_i(x_t)), E_{r^*} u(f^{r^*}(x_t)) \right\} + \right.$$

$$\left. - \underset{i \in \mathcal{F} \cup f^{r^*}}{\operatorname{argmax}} \left\{ \min_{\mu \in \mathcal{C}} E_{\mu,t-2} u(f_i(x_{t-1})), E_{r^*} u(f^{r^*}(x_{t-1})) \right\} \right\| > \epsilon \text{ infinitely often}$$

$$\Leftrightarrow \| \min_{\mu \in \mathcal{C}} E_{\mu,t-1} u(f) - \min_{\mu \in \mathcal{C}} E_{\mu,t-2} u(f) \| > \epsilon \text{ infinitely often}$$

$\Leftrightarrow^{\text{u is continuous and non constant}}$ ambiguity does not fade away. A contradiction. $\square$

References

Aumann, R. J. (1962). Utility theory without the completeness axiom. *Econometrica: Journal of the Econometric Society*, pages 445–462.

Avramov, D. (2002). Stock return predictability and model uncertainty. *Journal of Financial Economics*, 64(3):423–458.

Bewley, T. F. (2002). Knightian decision theory. part i. *Decisions in economics and finance*, 25(2):79–110.

Brown, G. W. (1951). Iterative solution of games by fictitious play. *Activity analysis of production and allocation*, 13(1):374–376.

Cremers, K. M. (2002). Stock return predictability: A bayesian model selection perspective. *The Review of Financial Studies*, 15(4):1223–1249.

De Rooij, S., Van Erven, T., Grünwald, P. D., and Koolen, W. M. (2014). Follow the leader if you can, hedge if you must. *The Journal of Machine Learning Research*, 15(1):1281–1316.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The quarterly journal of economics*, pages 643–669.

Epstein, L. G. and Schneider, M. (2007). Learning under ambiguity. *The Review of Economic Studies*, 74(4):1275–1303.

Epstein, L. G. and Seo, K. (2015). Exchangeable capacities, parameters and incomplete theories. *Journal of Economic Theory*, 157:879–917.

Freund, Y. and Schapire, R. E. (1999). Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2):79–103.

Fudenberg, D. and Levine, D. (1998). Learning in games. *European economic review*, 42(3):631–639.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of mathematical economics*, 18(2):141–153.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892.

Maccheroni, F., Marinacci, M., and Rustichini, A. (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74(6):1447–1498.

Marinacci, M. (2002). Learning from ambiguous urns. *Statistical Papers*, 43(1):143–151.

Marinacci, M. and Massari, F. (2017). Learning from ambiguous and misspecified models. *Working paper*.

Raiffa, H. (1961). Risk, ambiguity, and the savage axioms: comment. *The Quarterly Journal of Economics*, 75(4):690–694.

Savage, L. J. (1954). *The Foundations of Statistics*. Courier Corporation.

Timmermann, A. (2006). Forecast combinations. *Handbook of economic forecasting*, 1:135–196.

Walley, P. (1991). *Statistical reasoning with imprecise probabilities*. Chapman & Hall.