Ambiguity Aversion for Dummies

Daniel Méhkeri
dmehkeri@gmail.com

14 Jan 2012
DRAFT

Abstract

An ambiguity averse decision method is presented, and it is shown that it is dynamically consistent. Beliefs about events are represented as N-component vectors of probabilities. There is an extra N-component vector, the non-consequentialist ambiguity index. This is a special case of variational preferences [5]. Some of this paper is written for the decision theorist but the method itself is simple enough that it ought to be explainable at an undergraduate level, or an enriched high-school level, or perhaps even to an engineer.

1 The method

There is a simple, discrete version of the multiple-priors method of Gilboa and Schmeidler [1]: beliefs about events are represented as N-component vectors of probabilities. Each component of these vectors is separately additive, and they each update separately according to Bayes' rule. Known probabilities are represented by constant vectors.

This method is not dynamically consistent (unless $N = 1$). One way to restore dynamic consistency is to move to a discrete version of the variational preferences of Maccheroni, Marinacci, and Rustichini [5]. In fact the following method can be said to extend the simple N-component vector method in the same way that variational preferences extend multiple-priors.

- Beliefs about events are still represented as N-component vectors of probabilities. Each component is separately additive and updates according to Bayes' rule.
- There is an extra, *grounded* N-component vector, $\mathcal{A}$, the ambiguity index. *Grounded* means that the minimum component of this vector is zero.
- The (non-additive) measure of an event (or a binary act) is computed as the minimum component of the sum of its probability vector with the ambiguity index vector, $\min_i(P(f)_i + \mathcal{A}_i)$.
- To update $\mathcal{A}$, add to it the vector $P(f) - P(f|X)$, where f is the chosen (binary) act, and X is the updating event; also add a constant offset to each component so that $\mathcal{A}$ remains grounded, specifically $\max_i(P(f|X)_i - P(f)_i)$
- Acts with more than two outcomes are handled with a utility function, as usual.

The ambiguity index is non-consequentialist, in that there is a dependency on the counterfactual component of the chosen act. However there is no dependence on the feasible set.

This can also be generalised so that the number of components need not be specified a priori, without sacrificing explainability. Beliefs are represented by periodic sequences of probabilities. This can be explained as having vectors of different lengths. When necessary, two vectors can be padded out by repetition to the least common multiple of their lengths.

2 Example

Let us recall Machina's example of a mother with one indivisible treat and two children [4]. She strictly prefers to toss a coin to see who gets the treat, rather than giving the treat outright to either child. However, after the coin is tossed, she strictly prefers to give the treat to the winner, rather than toss another coin.

The analogous dynamically consistent and ambiguity averse behaviour occurs with the method of section 1. Let us consider the two-urn Ellsberg setup. We consider one draw from each urn. X and $\neg X$ are the two possible results of the draw from the clear urn, so that $P(X) = P(\neg X) = (\frac{1}{2}, \frac{1}{2})$. Y and $\neg Y$ are the two possible results of the draw from the unclear urn. Let us say the agent assigns prior probabilities $P(Y) = (\frac{1}{3}, \frac{2}{3}), P(\neg Y) = (\frac{2}{3}, \frac{1}{3})$. X and Y are independent so $P(X \wedge Y) = (\frac{1}{6}, \frac{1}{3})$ and so on.

Now consider the randomised event $Z = (X \wedge \neg Y) \vee (Y \wedge \neg X)$. It follows that $P(Z) = (\frac{1}{2}, \frac{1}{2})$, and its measure is $\frac{1}{2}$, therefore a bet on Z is strictly preferred to a pure bet on Y or $\neg Y$, both of which have measure $\frac{1}{3}$. Suppose the agent has chosen to bet on Z, and then the draw from the first urn occurs – say, the result is X – while the draw from the second urn has not yet occurred. The probability vector $P(Y|X) = P(Y)$ because they are independent. But Z has become a bet on $\neg Y$; $P(Z|X) = (\frac{2}{3}, \frac{1}{3})$. The ambiguity index updates to $\mathcal{A}^*_i = P(Z)_i - P(Z|X)_i + \max_{i'}(P(Z|X)_{i'} - P(Z)_{i'}) = (0, \frac{1}{3})$. Thus, the sum $P(Z|X) + \mathcal{A}^* = (\frac{2}{3}, \frac{2}{3})$, so the measure of Z will be $\frac{2}{3}$. Whereas, if we consider a new draw from the first urn with possible results X' and $\neg X'$ and a new randomised bet $Z' = (X' \wedge \neg Y) \vee (Y \wedge \neg X')$, then $P(Z') + \mathcal{A}^* = (\frac{1}{2}, \frac{5}{6})$ so its measure will still be $\frac{1}{2}$.

So, the agent is indifferent between a direct bet on either colour of the unclear urn, and strictly prefers a bet randomised on a draw from the clear urn, but, after the draw from the clear urn occurs, the agent strictly prefers to keep the bet on the chosen colour for the unclear urn, rather than switch to a bet randomised on another draw from the clear urn. This is analogous to the mother's preferences.

This simple example already captures a lot of reasonability conditions. It is supposed to be better to bet on the clear urn than the unclear urn, so there must be ambiguity aversion. The randomised event has objective probability $\frac{1}{2}$, so it must also be assigned that probability. However the result of a draw from the clear urn must not change the prior beliefs about the future draw from the unclear urn. If the agent acts only on these beliefs, in a consequentialist manner, then it would prefer to draw again from the clear urn, which is dynamically inconsistent. Of the methods that are not already ruled out, the method of section 1 (with $N = 2$) seems to me to be the simplest.

3 Dynamic Consistency

Theorem 3.1 *The method of section 1 is dynamically consistent, in the following strong sense: if f is the chosen act, and if neither X nor $\neg X$ are null events, and if g is weakly preferred to f after updating on a X, and also weakly preferred to f after updating on $\neg X$, then, g is weakly preferred to f before updating; moreover, if one of the posterior preferences is strict, then the prior preference is also strict.*

Proof: Let $\mathcal{A}$ be the prior ambiguity index vector. Let us write $\mathcal{A}(X;f)$ for the posterior given X, and $\mathcal{A}(\neg X;f)$ for the posterior given $\neg X$ when f is the chosen act. Let C^+ and C^- be the constant offsets required to keep $\mathcal{A}(X;f)$ and $\mathcal{A}(\neg X;f)$ grounded, respectively. Thus:

$$\begin{aligned}
\mathcal{A}(X;f)_i &= \mathcal{A}_i + P(f)_i - P(f|X)_i + C^+ \\
\mathcal{A}(\neg X;f)_i &= \mathcal{A}_i + P(f)_i - P(f|\neg X)_i + C^- \\
C^+ &= \max_i (P(f|X)_i - P(f)_i) \\
C^- &= \max_i (P(f|\neg X)_i - P(f)_i)
\end{aligned}$$

Let us write $V(f)$ and $V(g)$ for the prior measures of f and g, and $V(f|X;f)$, $V(g|X;f)$ for the posterior measures given X. By hypothesis g is weakly preferred to the chosen act f after updating on X, so:

$$V(g|X;f) \geq V(f|X;f)$$

So,

$$\min_i (P(g|X)_i + \mathcal{A}(X;f)_i) \geq \min_i (P(f|X)_i + \mathcal{A}(X;f)_i)$$

So,

$$\min_i (P(g|X)_i + \mathcal{A}_i + P(f)_i - P(f|X)_i + C^+) \geq \min_i (P(f|X)_i + \mathcal{A}_i + P(f)_i - P(f|X)_i + C^+)$$

So,

$$\min_i (P(g|X)_i + \mathcal{A}_i + P(f)_i - P(f|X)_i) \geq \min_i (\mathcal{A}_i + P(f)_i)$$

So, for any j,

$$P(g|X)_j + \mathcal{A}_j + P(f)_j - P(f|X)_j \geq \min_i (\mathcal{A}_i + P(f)_i)$$

And by similar reasoning for $\neg X$

$$P(g|\neg X)_j + \mathcal{A}_j + P(f)_j - P(f|\neg X)_j \geq \min_i (\mathcal{A}_i + P(f)_i)$$

So,

$$\begin{aligned}
&P(X)_j[P(g|X)_j + \mathcal{A}_j + P(f)_j - P(f|X)_j] \\
&+P(\neg X)_j[P(g|\neg X)_j + \mathcal{A}_j + P(f)_j - P(f|\neg X)_j] \\
&\geq P(X)_j[\min_i (\mathcal{A}_i + P(f)_i)] + P(\neg X)_j[\min_i (\mathcal{A}_i + P(f)_i)]
\end{aligned}$$

And of course $P(X)_j + P(\neg X)_j = 1$, so,

$$\begin{aligned}
&[P(X)_j P(g|X)_j + P(\neg X)_j P(g|\neg X)_j] \\
&+\mathcal{A}_j + P(f)_j \\
&-[P(X)_j P(f|X)_j + P(\neg X)_j P(f|\neg X)_j] \\
&\geq \min_i (\mathcal{A}_i + P(f)_i)
\end{aligned}$$

But the bracketed quantities are the prior probabilities, so,

$$P(g)_j + \mathcal{A}_j \geq \min_i (\mathcal{A}_i + P(f)_i)$$

And this is true for any j, that is,

$$\min_i(P(g)_i + \mathcal{A}_i) \geq \min_i(\mathcal{A}_i + P(f)_i)$$

So finally,

$$V(g) \geq V(f)$$

Which means g is weakly preferred to f before updating, as required. And the same proof works with one of the posterior preferences strict to show the prior preference is also strict. ■

Corollary 3.2 *The method of section 1 is dynamically consistent, in one of the senses discussed by Hanany and Klibanoff [2]: if f is the chosen act, and f is preferred (strictly or weakly) to g before updating on a non-null event X, and f and g agree on $\neg X$, then f is preferred (strictly or weakly, resp.) to g after updating.*

4 Interpretation

In addition to the method being simple to explain and dynamically consistent, it is desirable to be able to say what it means. The probability vectors by themselves just fall under multiple-priors [1] and are fairly intuitive. But we want to know what to make of the non-consequentialist ambiguity index vector, $\mathcal{A}$.

The measure of an act f is $V(f) = \min_i(P(f)_i + \mathcal{A}_i)$. When this measure is subtracted from this vector sum, what's left is the grounded part, $\mathcal{G}_i = P(f)_i + \mathcal{A}_i - V(f)$. So in the example of section 2, Y has measure $\frac{1}{3}$ and grounded part $(0, \frac{1}{3})$, $\neg Y$ has the same measure and grounded part $(\frac{1}{3}, 0)$, while the randomised event Z has measure $\frac{1}{2}$ and grounded part $(0, 0)$. After updating, $\neg Y$, therefore also Z, now has measure $\frac{2}{3}$ and grounded part $(0, 0)$, while the new randomised event Z' has measure $\frac{1}{2}$ and grounded part $(0, \frac{1}{3})$.

If f is the chosen act, then on update, its value becomes

$$\begin{aligned}
V(f|X; f) &= \min_i(P(f|X)_i + \mathcal{A}(X; f)_i) \\
&= \min_i(P(f|X)_i + \mathcal{A}_i + P(f)_i - P(f|X)_i) + \max_i(P(f|X)_i - P(f)_i) \\
&= \min_i(P(f)_i + \mathcal{A}_i) + \max_i(P(f|X)_i - P(f)_i) \\
&= V(f) + \max_i(P(f|X)_i - P(f)_i)
\end{aligned}$$

While the grounded part of the chosen act becomes

$$\begin{aligned}
\mathcal{G}^*_i &= P(f|X)_i + \mathcal{A}(X; f)_i - V(f|X; f) \\
&= P(f|X)_i + \mathcal{A}_i + P(f)_i - P(f|X)_i + \max_{i'}(P(f|X)_{i'} - P(f)_{i'}) - V(f|X; f) \\
&= P(f)_i + \mathcal{A}_i + \max_{i'}(P(f|X)_{i'} - P(f)_{i'}) - V(f|X; f) \\
&= P(f)_i + \mathcal{A}_i - V(f) \\
&= \mathcal{G}_i
\end{aligned}$$

So, $\mathcal{G}$ doesn't update on external events. It is a function only of the agent's choices. $\mathcal{A}$ is the part of $\mathcal{G}$ that resulted from choices that have become counterfactual. We can borrow a phrase from Bishop Berkeley and call $\mathcal{A}$ the ghost of a departed act.

Let us refer back to Machina's example. The mother wants only what is good for her children, and this involves *fairness*. If she gives the treat outright to either child, the good is diminished somewhat by her unfairness. It is fair that she decides to toss a coin, and this fairness remains after the coin is actually tossed.

In the context of ambiguous events, it is not totally clear what to call the analog of fairness among siblings, but we might tentatively call it *objectivity.* When $\mathcal{G}$ is zero the agent is acting objectively, and generally $\mathcal{G}$ tries to quantify the subjective nature of the agent's chosen act. This relates it back to the basic idea behind ambiguity averse behaviour in the first place.

References

[1] I. Gilboa and D. Schmeidler. *Maxmin expected utility with a non-unique prior.* J. Math. Econ. 18 (1989), pp 141-153.

[2] E. Hanany and P. Klibanoff. *Updating Ambiguity Averse Preferences.* B.E. Journal of Theoretical Economics (Advances) 2009 #9(1), article 37.

[3] E. Hanany, P. Klibanoff, and E. Marom. *Dynamically Consistent Updating of Multiple Prior Beliefs – an Algorithmic Approach.* Int. J. of Approximate Reasoning 52 (2011) pp 1198-1214.

[4] M. Machina. *Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty.* J. Econ. Lit. 27 (1989) pp 1622-1668.

[5] F. Maccheroni, M. Marinacci, and A. Rustichini. *Ambiguity aversion, robustness, and the variational representation of preferences.* Econometrica 74 (2006) pp 1447-1498.

[6] F. Maccheroni, M. Marinacci, and A. Rustichini. *Dynamic variational preferences.* J. Econ. Theory 128 (2006), pp 4-44.