Final accepted version of:

Klibanoff, Peter. Thoughts on Policy Choice under Ambiguity. *Journal of Institutional and Theoretical Economics* 169 (2013), 134-138. DOI: 10.1628/093245613X660492

Thoughts on Policy Choice under Ambiguity

Comment

by

Peter Klibanoff*

Manski (2013) presents an approach to the interesting and important problem of treatment allocation under ambiguity. In particular, there are a fixed number of mutually exclusive possible treatments, and a treatment allocation specifies what fraction of the population will be assigned to each. There is heterogeneity across the population in how individuals respond to each treatment. The ambiguity at issue comes from uncertainty about the population distribution of these treatment responses. This uncertainty is referred to as ambiguity, in the spirit of the literature following Ellsberg (1961), to emphasize that those deciding on the treatment allocation do not wish to treat this uncertainty in the same way that they would treat a known and well-understood risk. The first part of the paper assumes the choice of treatment allocation is made by a social planner, while the remainder considers the allocation being selected through voting or bilateral negotiation. The latter part includes consideration of the effect of the voting or negotiation rules being biased toward a status quo treatment. My comments directly address only the first part of the paper. This is both for the sake of brevity, and because the approach taken there, as well as my remarks, carries over in large part to the remainder.

Much of Manski's focus is on the desirability of treatment diversification – not assigning the whole population to the same treatment – as an optimal response to ambiguity. I think this focus is well placed, and agree that diversification is often a sensible course of action in such situations. Notice that diversifying across treatments can be seen as a way to (partially) hedge against the ambiguity by reducing the sensitivity of expected social welfare to the population distribution of treatment responses. The economic-theory and decision-theory literature concerned with ambiguity has had a long-standing emphasis on the fact that such hedging behavior may be valuable. In fact, the seminal papers of Schmeidler (1989) and Gilboa and Schmeidler (1989) *define* ambiguity aversion (they called it uncertainty aversion) as a preference for hedging. Klibanoff (2001) shows how the circumstances under which a strict preference for hedging is permitted can be used to characterize and distinguish between different models of ambiguity-averse preferences.

*Kellogg School of Management, Northwestern University. I thank Chuck Manski, Eran Hanany, and the seminar participants for helpful discussion, and the organizers for inviting me to comment.

To demonstrate the optimality of diversification from a social planner's perspective, Manski considers the case of two possible treatments, a and b. The planner knows only that the population distribution of treatment responses lies in some set $\mathcal{P}$. Corresponding to each distribution of treatment responses, there is an average population utility from treatment a and an average population utility from treatment b. Assume that the ambiguity makes the problem nontrivial in the sense that there is some distribution in $\mathcal{P}$ that makes expected social welfare under a better than under b and some for which the reverse is true. Manski endows the planner with a minimax-regret objective function. Thus the planner chooses the fraction, δ, of the population that will be assigned to treatment b as if, for each distribution $P \in \mathcal{P}$, she calculates a *regret* as the difference in expected social welfare between the best treatment given P and the allocation δ given P, and then chooses the δ that results in the lowest maximum (over $\mathcal{P}$) regret. He shows that whenever the problem is nontrivial, the minimax-regret solution involves some diversification. One criticism of this argument for diversification is that the minimax-regret criterion violates some normatively appealing choice properties. Notably, it generally leads choices to depend on unchosen alternatives in the feasible set (i.e., violates independence of irrelevant alternatives (IIA)), implying that it is inconsistent with any fixed preference ordering over treatment allocations. The following example serves to illustrate this and to suggest some alternative models.

Treatment Choice Example. Suppose the planner knows that there are only two possible population distributions of treatment responses, so that $\mathcal{P} = \{P_1, P_2\}$, and that treatment a is one whose expected treatment response is known (think of it as a status quo, with which there has been long experience), whereas b is a newly proposed treatment whose expected response is sensitive to the uncertainty about P. Specifically, let the expected social welfare for each treatment under each distribution in $\mathcal{P}$ be as follows:

	P_1	P_2
a	2	2
b	1	4

Then one can calculate that the expected social welfare for any treatment allocation $\delta \in [0, 1]$ is

	P_1	P_2
$\delta b + (1-\delta)a$	$2-\delta$	$2+2\delta$

From this, one can calculate the corresponding regrets:

	P_1	P_2
a	0	2
b	1	0
$\delta b + (1-\delta)a$	δ	$2-2\delta$

The treatment allocation that minimizes the maximal regret is therefore $\delta = 2/3$, i.e., assign $2/3$ of the population to the innovation, b, and $1/3$ to the status quo, a.

Now consider adding an additional innovative treatment option, c, that has expected social welfare a bit better than the status quo under P_1 but much less under P_2:

	P_1	P_2
c	2.1	0

The regrets relevant for calculating the minimax-regret treatment allocation can now be calculated as

	P_1	P_2
a	0.1	2
b	1.1	0
c	0	4
$\delta_b b + \delta_c c + (1 - \delta_b - \delta_c)a$	$0.1 + \delta_b - 0.1\delta_c$	$2 - 2\delta_b + 2\delta_c$

The treatment allocation that minimizes maximal regret is $\delta_b = 19/30$, $\delta_c = 0$, i.e., assign 19/30 of the population to b, 11/30 to a, and no one to c. Thus, this treatment allocation, which in the original problem was judged inferior to a $2/3, 1/3$ split between b and a, is now judged superior simply because an unused option c is present. Similar examples can be constructed where, instead of the addition of new unused options, it is the removal of some unchosen options from the original set that changes what is judged optimal.

Fortunately there alternatives to minimax regret that do satisfy IIA and are consistent with a preference ordering, yet still often recommend diversification in this and similar treatment-choice examples. They can be found among the models of ambiguity-averse preferences developed in the decision-theory literature. One model that fits the bill is the smooth ambiguity model (Klibanoff, Marinacci, and Mukerji, 2005). In the treatment-choice context, if $W_P(\delta)$ denotes the expected social welfare when the treatment allocation is δ and the population distribution of treatment responses is P, then the smooth ambiguity model evaluates δ by

$$E_\mu[\phi(W_P(\delta))],$$

where μ is a probability measure over the distributions in $\mathcal{P}$, E_μ means to take the expectation with respect to μ, and ϕ is an increasing function reflecting ambiguity attitude. When ϕ is concave, these preferences are ambiguity-averse. In the example, if $\phi(x) = (x^{1-\alpha})/(1-\alpha)$ with $0 < \alpha \neq 1$ and μ puts equal weight on P_1 and P_2, then it is strictly optimal to assign $\delta = 2(2^{1/\alpha} - 1)/(2 + 2^{1/\alpha})$ to treatment b when this expression is between 0 and 1, and the remainder to the status quo treatment a. This is true whether or not c is available. Notice that as long as α, the coefficient of relative ambiguity aversion, is greater than 1/2 , strict diversification is optimal; e.g., when $\alpha = 3/2$, $\delta \approx 0.33$.

More generally, there is a key property, beyond ambiguity aversion, that is essential in allowing the smooth ambiguity model to support diversification between a and b. It is the ability to depart from the certainty independence (or C-independence) axiom proposed in Gilboa and Schmeidler (1989). Klibanoff, Marinacci, and Mukerji (2005) view certainty independence as imposing constancy of *both* absolute and relative ambiguity attitude. The smooth ambiguity model exhibits this property only when ϕ is linear, reducing the model to expected utility. Functionally, certainty independence corresponds to the property of scale and translation invariance of indifference curves in utility space (Ghirardato, Maccheroni, and Marinacci, 2004). Many of the earliest models in the ambiguity literature, such as the max-min expected-utility model (Gilboa and Schmeidler, 1989) and the Choquet expected-utility model (Schmeidler, 1989), require certainty independence. Thus, those models cannot be used to make the case for diversification involving an unambiguous treatment, like a, that has the same expected welfare across all $P \in \mathcal{P}$.

Dynamics. Manski also considers a multiperiod setting where in each period the planner chooses treatments for the current generation, generations live for one period, and there is learning in the sense that observations of outcomes from earlier generations may inform the treatment choice for later generations. The criterion he proposes in this setting is adaptive minimax regret (AMR): each period, the planner solves the myopic minimax-regret problem for the current generation, using all information available at that time. Therefore, compared to the one-shot setting, the distributions in $\mathcal{P}$ may be transformed through learning, but otherwise the problems in each period are identical. Since AMR is myopic, no weight is given to the value of social learning for improving future policies when choosing in any given period. To take this into account, one needs a forward-looking model and to consider trade-offs across generations due to social learning on the one hand and discounting the future on the other. It is worth remarking that for forward-looking models under ambiguity there are often violations of dynamic consistency if the learning is done by simply applying distribution-by-distribution Bayesian updating to the set $\mathcal{P}$ and then reapplying the decision criterion using the updated set. In other words, preferences updated in this way will result in unwillingness to carry out ex ante optimal information-contingent plans. This observation applies not just to minimax regret, but to essentially all ambiguity models that allow updating in response to learning any nonnull event and use a set of distributions in this way to generate, within each choice problem, a complete preference ordering. This observation suggests that simple Bayesian updating may not be the optimal way to learn in such situations. For a theory of optimal updating under ambiguity that attains dynamic consistency by explicitly linking the process of updating to the problem context, see Hanany and Klibanoff (2007, 2009). Hanany, Klibanoff, and Marom (2011) provide an algorithmic implementation of such rules, and Baliga, Hanany, and Klibanoff (2012) contains an application.

This brings me to a final point. One of the truly nice things about Manski's line of research has been his attention to the explicit generation of ambiguity through econometric/statistical identification problems. This creates a channel linking the end result of statistical inference to the ambiguity appearing in the choice problem. The approach to dynamically consistent updating mentioned above suggests that it may also be important to consider the link in the reverse direction – that the features of the choice problem faced, such as the available treatment options, may optimally influence the manner in which statistical inference from new data is carried out.

References

Baliga, S., E. Hanany, and P. Klibanoff (2012), "Polarization and Ambiguity," http://www.kellogg.northwestern.edu/faculty/klibanof/ftp/workpap/Polarization and Ambiguity.pdf, retrieved July 31, 2012.

Ellsberg, D. (1961), "Risk, Ambiguity and the Savage Axioms," *The Quarterly Journal of Economics*, 75(4), 643–669.

Ghirardato, P., F. Maccheroni, and M. Marinacci (2004), "Differentiating Ambiguity and Ambiguity Attitude," *Journal of Economic Theory*, 118(2), 133–173.

Gilboa, I., and D. Schmeidler (1989), "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18(2), 141–153.

Hanany, E., and P. Klibanoff (2007), "Updating Preferences with Multiple Priors," *Theoretical Economics*, 2(3), 261–298.

— and — (2009), "Updating Ambiguity Averse Preferences," *The B.E. Journal of Theoretical Economics*, 9(1: Advances), doi: 10.2202/1935-1704.1547.

—, —, and E. Marom (2011), "Dynamically Consistent Updating of Multiple Priors Beliefs – An Algorithmic Approach," *International Journal of Approximate Reasoning*, 52(8), 1198–1214.

Klibanoff, P. (2001), "Characterizing Uncertainty Aversion through Preference for Mixtures," *Social Choice and Welfare*, 18(2), 289–301.

—, M. Marinacci, and S. Mukerji (2005), "A Smooth Model of Decision Making under Ambiguity," *Econometrica*, 73(6), 1849–1892.

Manski, C. F. (2013), "Status Quo Deference and Policy Choice under Ambiguity," *Journal of Institutional and Theoretical Economics (JITE)*, 169(1), forthcoming.

Schmeidler, D. (1989), "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57(3), 571–587.

Peter Klibanoff
Department of Managerial Economics
and Decision Sciences (MEDS)
Kellogg School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
U.S.A.
peterk@kellogg.northwestern.edu