Career Concerns and Ambiguity Aversion

March 9,2010 [Economics Letters, 108(2): 175–177 (August 2010)]

Eric Rasmusen

Abstract

Why do people have ambiguity aversion, preferring, a gamble with a 50% chance of success to one whose expected probability of success is 50% but where that 50% is an unbiased estimate? The answer modelled here, in the spirit of the career concerns literature, is learning: a risk-averse person does not wish observers to learn whether he is good or bad at estimating probabilities. He therefore prefers a gamble with objective probabilities.

Dan R. and Catherine M. Dalton Professor, Department of Business Economics and Public Policy, Kelley School of Business, Indiana University. BU 438, 1309 E. 10th Street, Bloomington, Indiana, 47405-1701. Office: (812) 855-9219. Fax:812-855-3354. erasmuse@indiana.edu, `http://www.rasmusen.org`. This paper: `http://www.rasmusen.org/papers/ambiguity-rasmusen.pdf`.

Keywords and Codes: D81, G31, L21, ambiguity, Ellsberg paradox, career concerns

I thank Nuffield College, Oxford University for its hospitality and Ashley Johnson for research assistance.

Introduction

Under our standard theory of expected utility, only expected utility matters, not how its probabilities are computed. A person should be indifferent between a gamble with a known, objective, probability of success to one with the same expected probability but greater dispersion. If he prefers the known probability we say he exhibits "ambiguity aversion." One variant of ambiguity aversion is to prefer a gamble with a single probability of success (e.g., .7) to one with a compound probability (e.g., 50% chance of .6, 50% chance of .8). The Ellsberg Paradox is an example (Ellsberg, 1961). Most people prefer a gamble in which they win if a red ball is drawn from urn A with 100 balls, 50 of which are red and 50 blue; to the same gamble with urn B which has 100 balls of an unknown color mix. It cannot be that they guess that there are less than 50 red balls in the second urn, because they also prefer urn A if they win when a blue ball is drawn. A second variant is to dislike making decisions that rely on estimated probabilities, as opposed to well-known, "objective" probabilities.

Most of the literature on ambiguity aversion (or "uncertainty aversion") has tried to define and axiomatize it, e. g. Schmeidler (1989), or, for a recent representative, Klibanoff, Marinacci and Mukerji (2005). Yet some cases of ambiguity aversion would seem to fit into the usual framework of risk aversion, allowing "strategy-based explanations" of the kind used in Harbaugh (forthcoming) to mimic the predictions of prospect theory. Halevy and Feltkamp (2005) explain it as the result of the possibility of positive correlation among estimate errors when there are bundled risks or Morris (1997), which bases it on rule rationality and the distrust of experimenters or other strategic actors.

I, too, will also use a strategy-based explanation to explain ambiguity aversion. The approach is in the spirit of the career concerns literature, in which agents take actions not just to maximize current performance, but to protect their reputation for high ability. Career concerns will make the agent averse to ambiguity if he is already averse to risk. Agents dislike ambiguity because it can make them look stupid. If their utility is concave in reputation, they will avoid a choice which displays their ability. In most models of career concerns and method choice, what is to be learned is the agent's ability to influence the probability of success, e.g. Milbourn, Shockley and

Thakor (2001). Here, it will be the agent's ability to estimate the probability of success. An agent's ability will become better known if observers learn something about the agent's subjective probability of success and something about the objective probability of success.

In the model, a risk-averse agent will choose one of two methods to achieve some goal. An example would be a doctor choosing either an old medical technique of known risk or a new one whose probability of success is only slowly becoming known. The agent will prefer a known probability of success because choosing the ambiguous method will reveal information on his ability to estimate probabilities. This information affects his payoff either directly, through his self-esteem, or indirectly, through his future career. In either case, risk aversion makes him prefer not to learn his ability.

This story depends on the agent not knowing his own ability. If he knows he is high-ability, then he will actually be ambiguity-loving, wanting a chance to reveal his ability. This may be what drove the results in one well-known experiment: Heath and Tversky (1990) found that subjects actually preferred ambiguous choices when they felt knowledgeable about the probabilities being estimated.

The Model

An unambiguous method, A, is successful with probability p_a. An ambiguous method, B, is successful with probability p_b, which equals with equal probability either β or $\alpha > \beta$. Profit is 0 if the chosen method is a failure, 1 if a success.

The agent gets a private signal *Alpha* or *Beta* of method B's success probability. With probability .5 he is untalented ($T = 0$) and his signal is correct with probability $\theta = \underline{\theta}$. With probability .5 he is talented ($T = 1$) and his signal is correct with probability $\theta = \overline{\theta}$, where $.5 \leq \underline{\theta} < \overline{\theta} \leq 1$.

Neither market nor agent knows if he is talented. Denote the market's estimate of the probability that he is talented by $\widehat{T}$. Denote the agent's estimate of p_b by $\widehat{p}_b$.

The agent's payoff function is

$$u(profit) + v(\widehat{T}), \tag{1}$$

where $u', v' > 0$ and $u'', v'' < 0$.

The agent might care about $\widehat{T}$ simply from pride, or because his future compensation depends on whether the market thinks he is good at estimating probabilities. A talented agent will have higher expected profit than an untalented one, and so would be paid more if his talent were known.

The agent's expected payoffs from methods A and B are thus

$$[p_a u(1) + (1 - p_a)u(0)] + v(.5) \tag{2}$$

and

$$[\widehat{p_b} u(1) + (1 - \widehat{p_b})u(0)] + [\widehat{p_b} v(\widehat{T}|S) + (1 - \widehat{p_b})v(\widehat{T}|F)] \tag{3}$$

These assumptions imply that the agent is risk averse and would prefer a safe method with a profit of p to a gamble. Nonetheless, the ambiguity of method B's probability of success does not matter in the profit part of the agent's utility function. If p_a and $\widehat{p_b}$ both equal some number p, then the expected payoffs are

$$[pu(1) + (1 - p)u(0)] + v(.5) \tag{4}$$

and

$$[pu(1) + (1 - p)u(0)] + [pv(\widehat{T}|S) + (1 - p)v(\widehat{T}|F)] \tag{5}$$

Reputation

When the agent receives the signal *Alpha* or *Beta* he does not thereby learn anything about his ability. Both qualities of method B have equal probability, so he is equally likely to receive each signal.

The agent's expected probability of success for method B if his private signal is *Beta* is

$$\widehat{p_b} = Pr(S|Beta) \quad = .5(\underline{\theta}\beta + (1 - \underline{\theta})\alpha) + .5(\overline{\theta}\beta + (1 - \overline{\theta})\alpha) \tag{6}$$

because with probability .5 he is untalented, in which case the *Beta* signal is correct with probability $\underline{\theta}$ so $p_b = \beta$, and incorrect with probability $1 - \underline{\theta}$ so $p_b = \alpha$. With probability .5 he is talented, in which case the *Beta* signal is correct with probability $\overline{\theta}$ so $p_b = \beta$, and incorrect with probability $1 - \overline{\theta}$ so $p_b = \alpha$.

As a running example, let $\beta = .3, \alpha = .7, \underline{\theta} = .6, \overline{\theta} = .9$. Then

$$\widehat{p_b} = .5((.6)(.3) + (.4)(.7)) + .5((.9)(.3) + (.1)(.7)) = .40. \tag{7}$$

Similarly, if his private signal is *Alpha*,

$$\begin{aligned} \widehat{p_b} = Pr(S|Alpha) &= .5((1-\underline{\theta})\beta + \underline{\theta}\alpha) + .5((1-\overline{\theta})\beta + \overline{\theta}\alpha) \\ &= .5((.4)(.3) + (.6)(.7)) + .5((.1)(.3) + (.9)(.7)) = .60. \end{aligned} \tag{8}$$

If $p_a < E(p_b|Beta)$, the agent would adopt method B regardless of his signal. The market could not learn anything about his talent, regardless of whether it observed the success or failure of the method, or even the true probability p_b.

If method B is chosen and $E(p_b|Beta) < p_a < E(p_b|Alpha)$, the market can deduce that the signal was *Alpha*. This has no immediate use in learning the agent's talent, but in conjunction with observing S or F it will be useful.

After observing S or F, the market forms its posterior belief about the agent's talent. If the method is successful and the signal was *Alpha*, then using Bayes's Rule, the belief is

$$\begin{aligned} \widehat{T}(S) = Pr(T=1|S, Alpha) &= \frac{Pr(S|T=1, Alpha)Pr(T=1)}{Pr(S|Alpha)} \\ &= \frac{(.5(1-\overline{\theta})\beta + .5\overline{\theta}\alpha)(.5)}{(.5(1-\overline{\theta})\beta + .5\overline{\theta}\alpha)(.5) + (.5(1-\underline{\theta})\beta + .5\underline{\theta}\alpha)(.5)} \\ &= \frac{[(.5)(.1)(.3) + (.5)(.9)(.7)](.5)}{[(.5)(.1)(.3)+(.5)(.9)(.7)](.5) + [(.5)(.4)(.3)+(.5)(.6)(.7)](.5)} = .55. \end{aligned} \tag{9}$$

or

$$\widehat{T}(F) = Pr(T=1|F, Alpha) = \frac{Pr(F|T=1, Alpha)Pr(T=1)}{Pr(F|Alpha)}$$

$$= \frac{(.5\overline{\theta}(1-\alpha)+.5(1-\overline{\theta})(1-\beta))(.5)}{(.5\overline{\theta}(1-\alpha)+.5(1-\overline{\theta})(1-\beta))(.5)+(.5\underline{\theta}(1-\alpha)+.5(1-\underline{\theta})(1-\beta))(.5)}$$

$$= \frac{[.5(.9)(.3)+.5(.1)(.7)](.5)}{[.5(.9)(.3)+.5(.1)(.7)](.5)+[.5(.6)(.3)+.5(.4)(.7)](.5)} = .425. \quad (10)$$

The expected value of $\widehat{T}$ is

$$\begin{aligned} ET &= Pr(S|Alpha)Pr(T=1|S, Alpha) \\ &\quad + Pr(F|Alpha)Pr(T=1|F, Alpha) \\ &= \overline{\theta}(.5) + (1-\overline{\theta})(.5) = .5 \end{aligned} \quad (11)$$

where the second line is obtained by substituting from (9) and (10).

If the agent was risk-neutral in reputation, his decision to accept the method would be unaffected by the market learning about his ability, but we have assumed he is risk averse. By Jensen's Inequality, $E(v(\widehat{T})) < v(E(\widehat{T}))$ because v is concave and $\widehat{T}$ is stochastic. If $\widehat{p_b} = p_a$, the agent would choose method A, exhibiting ambiguity aversion.

Thus, we can explain ambiguity aversion as the result of career concerns without any need to resort to irrationality or direct disutility. The agent's principal should design compensation to encourage the agent to take on ambiguous methods: (1) to get him to choose the most profitable method, since otherwise he might prefer method A even if $p_a < \widehat{p_b}$; and (2) to learn his ability to decide whether to keep him or hire a new agent. This second advantage would make the principal actually ambiguity loving.

Fifty-Fifty Hindsight

"Hindsight bias" is the tendency of people to fool themselves into believing ex post that they would have made a different decision ex ante. If we increase the ex post information available to the market, the market will punish or reward the agent in a way that looks like hindsight bias but is actually unbiased.

In the last section's model, when the market observes S it deduces that the agent, having observed $Alpha$, was more likely than not talented and the true probability was α. Now assume that once the method is completed the market learns whether $p_b = \alpha$ directly— regardless of actual success, the ex ante probability of success becomes common knowledge. This is even better information for learning about the agent's ability.

If the market observes p_b directly, it does not need to use the imperfect signal of S or F. Rather, $\widehat{T}$ becomes

$$\begin{aligned} \widehat{T}(p_b = \alpha) &= Pr(T = 1|Alpha, p_b = \alpha) = \tfrac{Pr(p_b=\alpha|T=1, Alpha)Pr(T=1)}{Pr(p_b=\alpha|Alpha)} \\ &= \tfrac{\overline{\theta}(.5)}{\overline{\theta}(.5)+\underline{\theta}(.5)} = .6. \end{aligned} \tag{12}$$

or

$$\begin{aligned} \widehat{T}(p_b = \beta) &= Pr(T = 1|Alpha, p_b = \beta) = \tfrac{Pr(p_b=\beta|T=1, Alpha)Pr(talented)}{Pr(p_b=\beta|Alpha)} \\ &= \tfrac{(1-\overline{\theta})(.5)}{(1-\overline{\theta})(.5)+(1-\underline{\theta})(.5)} \\ &= \tfrac{(.1)(.5)}{(.1)(.5)+(.4)(.5)} = \tfrac{.05}{.05+.20} = .2. \end{aligned} \tag{13}$$

The gap between $\widehat{T}(p_b = \beta)$ and $\widehat{T}(p_b = \alpha)$ is bigger than the gap between $\widehat{T}(F)$ and $\widehat{T}(S)$ because

$$\widehat{T}(p_b = \alpha) = \frac{\overline{\theta}(.5)}{\overline{\theta}(.5) + \underline{\theta}(.5)} = \frac{\overline{\theta}\alpha}{\overline{\theta}\alpha + \underline{\theta}\alpha} > \frac{\overline{\theta}\alpha + W_1}{\overline{\theta}\alpha + W_1 + \underline{\theta}\alpha + W_2} \tag{14}$$

if $W_1 < W_2$. From equation (9), however, $\widehat{T}(S)$ equals that right-hand expression in (14), however, if we set $W_1 = (1 - \overline{\theta})\beta$ and $W_2 = (1 - \underline{\theta})\beta$, so $\widehat{T}(p_b = \alpha) > \widehat{T}(S)$ not just for our numerical example but in general. It can similarly be shown that $\widehat{T}(p_b = \beta) < \widehat{T}(F)$.

Thus, if the market has 50-50 hindsight and observes p_b directly instead of just observing S or F, the agent's risk from the ambiguous method rises. He has reason to prefer a method whose success probability is objectively known to the public rather than a method whose success probability he has estimated, even if the probabilities turn out the same. Furthermore, the agent is not just reacting to a behavioral bias on the part of the market: if what would have been the best decision ex ante becomes clear ex post, the market can use the information to better learn the agent's ability.

References

Ellsberg, D., 1961, Risk, ambiguity, and the Savage axioms, Quarterly Journal of Economics 75, 643–669.

Halevy, Y. and V. Feltkamp, 2005, A Bayesian approach to uncertainty aversion, Review of Economic Studies 72, 449–466.

Harbaugh, R., forthcoming, Prospect theory or skill signaling? American Economic Review.

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

Milbourn, T., R. Shockley and A. Thakor, 2001, Managerial career concerns and investments in information, RAND Journal of Economics 32, 334–351.

Morris, S., 1997, Risk, uncertainty and hidden information, Theory and Decision 42, 235–269.

Schmeidler, D.,1989, Subjective probability and expected utility without additivity, Econometrica 57, 571–587.