

SHARING AMBIGUITY

Larry G. Epstein*

In a version of the Ellsberg Paradox, the decision-maker is confronted with two urns, each containing 100 balls that are either Red or Blue. She is told that there are 50 of each color in the first ('unambiguous') urn, but no further information is provided about the second ('ambiguous') urn. There is a widely exhibited preference to bet on drawing Red (or Blue) from the first urn rather than from the second. Though such rankings are intuitive, they are inconsistent with subjective expected utility theory and, more generally, with reliance on *any* single probability measure to represent beliefs. Thus the Paradox illustrates the behavioral meaning of the Knightian distinction between risk (measurable or probabilistic uncertainty) and ambiguity (unmeasurable uncertainty).

The importance of the Ellsberg Paradox is the intuition that this distinction may be important much more widely. In particular, it seems at least plausible to view consumption-savings and portfolio choice decisions as being qualitatively different than the choice of which bet to accept on the outcome of a coin flip; only the latter is a choice between risky prospects. My objective in this paper is to illustrate both the tractability and potential fruitfulness (for addressing the home-bias puzzle, for example) of a macro-style model that permits aversion not only to risk but also to ambiguity.

I employ a simple two-period heterogeneous-agent economy. The time periods are $t - 1$ ('today') and t ('tomorrow'). Uncertainty is represented by the state space Ω . There are two consumers and consumer i 's consumption process is (c_{t-1}^i, c_t^i) , where c_{t-1}^i is deterministic, c_t^i is a random variable on Ω and each is real-valued. I consider an endowment economy with aggregate endowment (Y_{t-1}, Y_t) , where Y_t is random. The efficient allocation of this endowment is usually posed as a problem of efficient risk sharing. In particular, it is assumed that consumers' beliefs are represented by a common probability measure. If it is assumed further that vNM indices are (increasing, concave and) additive across time and that consumers have a common discount factor, then efficient allocations are such that each c_t^i is an increasing function of Y_t . Consequently, consumption is perfectly correlated across

consumers and, if preference is homothetic, then consumption growth rates are equal. These predictions are often contradicted dramatically by data, particularly in international settings where consumers represent countries and where individual country growth rates respond to idiosyncratic shocks. See Karen Lewis (1999) for a survey; she terms the observed systematic violation of efficient risk sharing the consumption home-bias puzzle.

The model that I outline continues to assume complete markets and hence focuses on efficient allocations. However, taking Ellsberg seriously, I drop the assumption that it is risk alone that is to be shared. I assume that consumers are not sufficiently confident to assign sharp probabilities to all future states. Rather, following Itzhak Gilboa and David Schmeidler (1989), beliefs are represented by a (nonsingleton) set of priors and consumption prospects are ranked according to their minimum expected utility as probabilities vary over the set of priors. Thus consumers view consumption prospects as ambiguous and the question of interest is “what is the nature of efficient sharing of ambiguity?” Though I will refer to consumers as countries, the international interpretation is evidently optional.¹

I. AN ECONOMY WITH AMBIGUITY

In order to accommodate idiosyncratic shocks for each of two countries, the state space is taken to be the two-dimensional set $\Omega = \{-1, 1\} \times \{-1, 1\}$ and the corresponding driving state process is $W_t = (W_t^1, W_t^2)$, where $W_t^i(\omega) = \omega^i$, $i = 1, 2$. The equal-probability measure P on Ω implies that each W_t^i has zero mean and unit variance and that W_t^1 and W_t^2 are independently distributed.

I assume that i is more familiar with her own (domestic) component process W_t^i than with the other (foreign) component W_t^j .² In extreme form this leads to no ambiguity for i about W_t^i , though W_t^j is ambiguous for her. In particular, while i assigns equal probabilities to the two possible outcomes of W_t^i , for the foreign process she is confident only that the probability of each possible outcome lies in the interval $[\frac{1-\kappa^i}{2}, \frac{1+\kappa^i}{2}]$, where $0 \leq \kappa^i < 1$ is a parameter representing the extent of her ambiguity about W_t^j . Domestic and foreign shocks are viewed as independent. Accordingly, 1’s set of priors on Ω , \mathcal{P}^1 , consists of all products

of the $(\frac{1}{2}, \frac{1}{2})$ measure on the first component space (for W_t^1) with measures on the second component space lying in the appropriate interval. Define \mathcal{P}^2 similarly. Thus each country faces an analogue of the 2-urn Ellsberg setting, though the identity of the ambiguous urn differs between countries consistent with the subjective nature of ambiguity. Note that each set of priors contains P .

The description of the economy is completed by specifying utilities and the endowment. Country i 's utility function is

$$V^i(c_{t-1}^i, c_t^i) = \log c_{t-1}^i + \beta \min_{Q^i \in \mathcal{P}^i} E_{Q^i} \log c_t^i, \quad (1)$$

where $0 < \beta < 1$. The standard logarithmic model is obtained in the special case $\kappa^i = 0$. The aggregate endowment is Y_{t-1} and Y_t , where

$$Y_t(\omega) / Y_{t-1} = \exp(\mu^Y + s^Y \cdot \omega) / E_P [\exp(s^Y \cdot W_t)] .$$

Thus $e^{\mu^Y} - 1$ is the expected growth rate of the endowment according to P . It is without loss of generality to assume that $s^Y \geq 0$, which normalization brands $W_t^i = 1$ as a good realization and $W_t^i = -1$ as a bad one. For concreteness, suppose that $s^Y > 0$.

II. EFFICIENT ALLOCATIONS

Efficient allocations solve the planning problem

$$\begin{aligned} \max V^1(c_{t-1}^1, c_t^1) + \lambda V^2(c_{t-1}^2, c_t^2) \\ \text{s.t. } c_\tau^1(\cdot) + c_\tau^2(\cdot) = Y_\tau(\cdot), \quad \tau = t-1, t, \end{aligned} \quad (2)$$

where $\lambda > 0$ is the relative utility weight for country 2. At any allocation and resulting consumption for 1, there is a measure Q^1 that solves the minimization in (1). Then Q^1 is completely described by the probability, denoted $\frac{1+\theta^1}{2}$, that it assigns to the event $W_t^2 = 1$. Similarly for country 2. Thus an envelope theorem implies the same first-order conditions as

would apply for a planning problem in which sets of priors are replaced by the single prior Q^i for $i = 1, 2$. That is, $c_{t-1}^2 = \lambda c_{t-1}^1$ and $c_t^2 = \lambda \rho_t c_t^1$, where

$$\rho_t(\omega) = (1 + \theta^2 \omega^1) / (1 + \theta^1 \omega^2). \quad (3)$$

Deduce that

$$\begin{aligned} c_{t-1}^1 &= \frac{1}{1 + \lambda} Y_{t-1}, & c_{t-1}^2 &= \frac{\lambda}{1 + \lambda} Y_{t-1}, \\ c_t^1 &= \frac{1}{1 + \lambda \rho_t} Y_t, & c_t^2 &= \frac{\lambda \rho_t}{1 + \lambda \rho_t} Y_t. \end{aligned} \quad (4)$$

These expressions do not fully describe efficient allocations because the θ^i 's and hence ρ_t are endogenous. Since θ^i corresponds to a subjectively worst measure in \mathcal{P}^i , one might expect that it equal an extreme point $\pm \kappa^i$. In fact, that is not necessarily true as indicated in the complete description of efficient allocations that follows.

Theorem: Write $\theta = (\theta^1, \theta^2)$ and $\kappa = (\kappa^1, \kappa^2)$. Define the functions $\Phi^i(\theta)$ for $-\kappa \leq \theta \leq \kappa$ by

$$\Phi(\theta) = \begin{bmatrix} \Phi^1(\theta) \\ \Phi^2(\theta) \end{bmatrix} = \begin{bmatrix} \int \log [1 + (\lambda \rho_t)^{-1}] \omega^1 dP \\ \int \log [1 + \lambda \rho_t] \omega^2 dP \end{bmatrix},$$

where ρ_t is defined in (3). Then an allocation solves (2) if and only if it has the form (4) where θ is the unique solution to

$$\Phi(\theta) \leq s^Y \text{ and } (\Phi(\theta) - s^Y) \cdot (\theta + \kappa) = 0. \quad (5)$$

The latter solution satisfies $-\kappa \leq \theta \leq 0$.

The formalism surrounding (5) suggests an interpretation whereby $\Phi(\cdot)$ represents the demand for volatility, s^Y is the supply and these are equilibrated by adjustment of θ , the relevant (constrained) 'price'. Consistent with this interpretation, we have a complementary slackness condition, whereby for each i ,

$$\text{either } [\Phi^i(\theta) < s_i^Y \ \& \ \theta^i = -\kappa^i] \text{ or } \Phi^i(\theta) = s_i^Y.$$

In the special case $\Phi(-\kappa) < s^Y$, then $\theta = -\kappa$, which means that each country acts as though she attaches the smallest possible probability to good realizations of the foreign

shock. Consequently, (4) represents a closed-form solution. Because $\Phi(0) = 0$, this case applies for small κ or large s^Y . But the Theorem covers also the case of large κ (small s^Y), where efficiency implies

$$\Phi(-\kappa) > s^Y \quad \text{and hence } \theta > -\kappa, \quad (6)$$

that is, countries do *not* act as though an extreme point in their set of priors applies. Here beliefs, in the sense of the shadow singleton prior for each agent, are selected endogenously in equilibrium.

Uniqueness of the solution θ implies uniqueness of the Arrow-Debreu prices that support the efficient allocation corresponding to any given λ . This is in contrast to my paper with Tan Wang (1994) that emphasizes the potential of ambiguity for generating price indeterminacy. It contrasts also with Truman Bewley (1998), who in his closely related model points to the fact that preferred sets in a 2-state Edgeworth box typically have corners, supporting his claim that Knightian uncertainty inhibits trade. The difference here is the special asymmetric structure of ambiguity whereby i is ambiguous only about W_t^j . Thus i 's probabilistic beliefs about her own process W_t^i pin down prices for consumption contingent on W_t^i states.

A number of qualitative properties of efficient allocations can be derived from the theorem.³ Consumption in each country is non-negatively correlated with shocks in *both* countries. For the domestic shock, this is evident from (4); for the foreign shock, (5) implies

$$E_P [W_t^j \log c_t^i] \geq 0, \quad i \neq j.$$

The extreme of equality with zero occurs precisely in allocations corresponding to (6); for example, if the ambiguity parameter κ is large. Thus even when i 's ambiguity about W_t^j is large, it is not efficient for her to 'short' the foreign shock.

A second implication is that consumption growth rates are not perfectly correlated across countries. In fact, idiosyncratic consumption growth rates are positively correlated with idiosyncratic shocks in the sense that (if $\kappa^j > 0$)

$$cov_P(\log(c_t^i/c_{t-1}^i) - \log(Y_t/Y_{t-1}), W_t^i) > 0. \quad (7)$$

Comparative statics analysis of (5) yields that (i) each θ^i is (weakly) decreasing in each s_j^Y and (ii) θ^2 is decreasing and θ^1 increasing in λ . Recall that $(1 + \theta^i)/2$ can be interpreted as the ambiguity-adjusted probability that i assigns, in equilibrium, to the good outcome $W_t^j = 1$. Accordingly, optimism in both countries declines with an increase in the volatility of aggregate consumption (due to an increase in s_i^Y or s_j^Y) and a redistribution towards country 2 (increase in λ) makes 2 more optimistic and 1 more pessimistic. Finally, if one measures the size of home-bias in each country by the covariance in (7), then redistribution towards country 2 reduces home-bias there and increases it in country 1.

III. PROOF OF THEOREM

I include a sketch of the nontrivial part of the proof in order to emphasize its simplicity and because it is informative also about the nature of arguments needed in a multi-period setting.

To show that every efficient allocation has the stated form, focus on the period t component of the planning problem (2), namely on

$$\begin{aligned} & \max_{(c_t^i)} \sum_i \lambda_i \min_{(Q^i \in \mathcal{P}^i)} E_{Q^i} u^i(c_t^i) = \max_{(c_t^i)} \min_{(Q^i)} \sum_i \lambda_i E_{Q^i} u^i(c_t^i) \\ & = \min_{(Q^i)} \max_{(c_t^i)} \sum_i \lambda_i E_{Q^i} u^i(c_t^i) = \min_{(|\theta^i| \leq \kappa^i)} \max_{(c_t^i)} \sum_i \lambda_i E_P[u^i(c_t^i) (1 + \theta^i W_t^j)] \\ & \equiv \min_{\theta} J(\theta; \lambda), \end{aligned}$$

where consumption levels are constrained by $\sum c_t^i \leq Y_t$, I have applied the minimax theorem to reverse the min and max operations, $j \neq i$ in the last summation and where

$$J(\theta, \lambda) \equiv \max_{(c_t^i)} \sum_i \lambda_i E_P[u^i(c_t^i) (1 + \theta^i W_t^j)].$$

The envelope theorem implies that (using the fact that each $u^i(\cdot) = \log(\cdot)$)

$$\begin{aligned} J_{\theta^1}(\theta, \lambda) &= E_P[W_t^2 u^1(c_t^1)] = E_P[W_t^2 (\log Y_t - \log(1 + \lambda \rho_t))] \\ &= s_2^Y - E_P[W_t^2 \log(1 + \lambda \rho_t)] = s_2^Y - \Phi^2(\theta), \end{aligned} \tag{8}$$

$$\begin{aligned}
J_{\theta^2}(\theta, \lambda) &= \lambda E_P[W_t^1 u^2(c_t^2)] = \lambda E_P[W_t^1 (\log Y_t - \log(1 + (\lambda\rho_t)^{-1}))] \\
&= \lambda s_1^Y - \lambda E_P[W_t^1 \log(1 + (\lambda\rho_t)^{-1})] = \lambda s_1^Y - \lambda \Phi^1(\theta).
\end{aligned}$$

Thus the Kuhn-Tucker Theorem implies (5).

The optimal θ must satisfy $-\kappa \leq \theta \leq 0$ because $-\kappa^i < \theta^i < \kappa^i \implies s_j^Y = \Phi^j(\theta^*)$, $i \neq j$, $\implies \theta^i < 0$. (By elementary arguments one can show that $\Phi^2(\theta^1, \theta^2) < 0$ if $\theta^1 > 0$ and $\Phi^1(\theta^1, \theta^2) < 0$ if $\theta^2 > 0$.) Similarly one can exclude an optimum at $+\kappa_1$: that would require $s_2^Y < \Phi^2(\kappa^1, \theta^2)$, but the latter is negative. This completes the proof.

The comparative statics analysis made use of the following: As a pointwise maximum of a collection of linear (in θ) functions, $J(\cdot, \lambda)$ is convex for each λ . Therefore, (8) implies that

$$-D_{\theta\theta}J(\theta, \lambda) = \begin{bmatrix} \partial\Phi^2/\partial\theta^1 & \partial\Phi^2/\partial\theta^2 \\ \lambda\partial\Phi^1/\partial\theta^1 & \lambda\partial\Phi^1/\partial\theta^2 \end{bmatrix}$$

and both matrices are negative definite. In particular, $\det(D_\theta\Phi(\theta)) < 0$ and $\partial\Phi^i/\partial\theta^j < 0$ for $i \neq j$.

IV. CONCLUDING REMARKS

The preceding model can be extended to a multi-period setting. Think of a two-dimensional state process $W_t = (W_t^1, W_t^2)$ that is a random walk under a reference probability measure. Suppose that while country i is confident that the domestic shock is a random walk, she views W_t^j as an *ambiguous random walk*, that is, conditional on the state at time $t - 1$, her beliefs are that $W_t^j - W_{t-1}^j = \pm 1$ according to the color of the ball drawn from an ambiguous Ellsberg urn. Thus conditional one-step-ahead beliefs have the same form as in the two-period model. Using them one can define utility recursively, essentially by replacing $\log c_t^i$ in (1) by $V_t^i(c^i)$, the continuation utility for periods t and beyond. The resulting model of single-agent utility admits the explicit representation

$$V_t^i(c^i) = \min_{Q \in \mathcal{P}^i} E_Q \left[\sum_{\tau \geq t} \beta^{\tau-t} u^i(c_\tau^i) \mid \mathcal{F}_t \right], \tag{9}$$

for a suitable set \mathcal{P}^i of priors over possible trajectories of W_t .⁴ This utility specification has a number of attractive features that I now describe.

First, it has a suitable continuous-time limit, as described in work with Werner Ploberger (2001), where the driving state process is an *ambiguous Brownian motion*. Jianjun Miao and I (2000) have applied the resulting model of utility to a two-country setting that is the continuous-time counterpart of this paper’s model. The analytical power of continuous-time permits sharp results to be derived; they confirm and extend those reported above. In particular, we describe the implementation of efficient allocations as a Radner equilibrium and describe asset market implications (home-bias in equities, for example) of ambiguity.

My paper with Martin Schneider (2001b) provides a simple axiomatic basis for a generalization of (9) in which \mathcal{P}^i is restricted to conform to the ‘spirit’ but not the letter of the above story about an ambiguous random walk. The essential characterizing axioms are: (i) each conditional utility V_t^i satisfies the axioms described by Gilboa and Schmeidler (1989) that characterize the multiple-priors model in an atemporal or one-shot choice framework; and (ii) the collection $\{V_t\}_{t \geq 0}$ of all conditional preferences is dynamically consistent. In contrast, the related model of utility proposed by Evan Anderson et al. (2000), which is also the subject of another paper in this session, violates dynamic consistency. In particular, plans that are optimal from the perspective of time 0 will generally not be carried out.⁵

Further, learning can be accommodated. Though the specific conditional one-step-ahead beliefs described above are the same at every node and thus do not respond to past observations, the model in its general axiomatic form permits such responsiveness to data (see my papers with Zengjing Chen (2000) and with Schneider (2001a)). Prior by prior application of Bayes’ Rule provides a dynamically consistent updating rule for recursive multiple-priors utility. Moreover, a rich set of learning dynamics is admitted. For example, in many environments ambiguity can plausibly persist indefinitely.⁶ In others, ambiguity may increase in response to a ‘surprising’ observation that leads the agent to doubt her previous view (model) of the environment.

An important outstanding question is “what are reasonable values for κ^i ?” One possible approach is to apply Bayesian detection theory for discriminating between probability laws in order to assess how difficult it would be to discriminate between measures lying in the

set of priors corresponding to a specific value for κ^i . This route has been developed by Anderson et al. (2000) for their model of robust decision-making; it seems likely that the approach could be adapted to our model. Alternatively, interpret the challenge as being the difficulty of transferring ambiguity parameters across settings. There is no difficulty transferring risk aversion parameters because any given lottery presumably represents the same prospect regardless of the context. In contrast, ambiguity is by its very nature tied to a specific state space. There is a need to uncover deeper structural parameters underlying the κ^i 's that are transferable across settings.

REFERENCES

Anderson, Evan; Hansen, Lars and Sargent, Tom. “Robustness, Detection and the Price of Risk.” Mimeo, Stanford University, 2000.

Bewley, Truman. “Knightian Uncertainty,” in Donald Jacobs, Ehud Kalai and M. Kamien, eds., Frontiers of Research in Economic Theory: the Nancy L. Schwartz Memorial Lectures. New York: Cambridge University Press, 1998, pp. 71-81.

Chateauneuf, Alain; Dana, Rose-Anne and Tallon, Jean-Marc. “Optimal Risk Sharing Rules and Equilibria with Choquet-Expected-Utility.” Journal of Mathematical Economics, 2000, (34), pp. 191-214.

Chen, Zengjing and Epstein, Larry G. “Ambiguity, Risk and Asset Returns in Continuous Time.” Mimeo, University of Rochester, 2000.

Epstein, Larry G. and Miao, Jianjun. “A Two-Person Dynamic Equilibrium.” Mimeo, University of Rochester, 2000.

Epstein, Larry G. and Ploberger, Werner. “Ambiguous Brownian Motion.” Mimeo, University of Rochester, 2001.

Epstein, Larry G. and Schneider, Martin. “Learning under Ambiguity.” Mimeo, University of Rochester, 2001a.

Epstein, Larry G. and Schneider, Martin. “Recursive Multiple-Priors Utility.” Mimeo, University of Rochester, 2001b.

Epstein, Larry G. and Wang, Tan. “Intertemporal Asset Pricing under Knightian Uncertainty.” Econometrica, March 1994, (62), pp. 283-322.

Gilboa, Itzhak and Schmeidler, David. “Maxmin Expected Utility with Nonunique Prior.” Journal of Mathematical Economics, 1989, (18), pp. 141-153.

Huberman, Gur. “Familiarity Breeds Investment.” Mimeo, Columbia University, 2000.

Lewis, Karen. “Trying to Explain Home Bias in Equities and Consumption.” Journal of Economic Literature, 1999, (37), pp. 571-608.

NOTES

* Department of Economics, U. Rochester, Rochester, NY 14627, lepn@troi.cc.rochester.edu.
The paper owes an obvious debt to Tan Wang, Zenjing Chen, Jianjun Miao, Werner Ploberger and Martin Schneider. This research was supported by the NSF (grant SES-9972442).

¹For a related paper dealing with the characterization of efficient allocations see, for example, Chateauneuf et al (2000). The two-period model that follows differs in that more concrete results are delivered as a result of strong functional form assumptions. More importantly, the model and its essential predictions may be extended to a multi-period dynamic setting as I describe below.

²See Gur Huberman (2000) for recent market evidence of the preference to bet on the familiar.

³These properties rely on the facts: (i) ρ_t is decreasing in W_t^1 and increasing in W_t^2 ; (ii) $cov(f(X), X) \geq 0$ for any random variable X and increasing function f ; and (iii) various properties of Φ that are available on request. The key property is described after the proof of the theorem.

⁴The model is a special case of that described by myself and Wang (1994).

⁵ The noted companion paper by Hansen et al disagrees. They claim in their discussion of their ‘constraint preferences’ model that it satisfies a weaker form of dynamic consistency but one that still rules out the revision of ex ante optimal plans. In my view, their notion of dynamic consistency is irrelevant to the issues at hand.

⁶Bewley’s (1998) discussion of learning under “Knightian uncertainty” is very relevant here.