Inference under Stability of Risk Preferences

Levon Barseghyan Francesca Molinari Joshua C. Teitelbaum

January 25, 2013

For an economic model to be useful, it is desirable that it has explanatory power that extends (at least) across contexts similar in nature. This idea can be formalized by requiring that the model exhibits "stability" — a single parameterization of the model should be consistent with behavior across different contexts. Recent work by Barseghyan, Prince, and Teitelbaum (2011) and Einav, Finkelstein, Pascu, and Cullen (2012), however, suggests that the standard model of risky choice — the expected utility (EU) model — fails to satisfy this stability criterion. Both studies find that, under a standard EU model, a majority of households (more than two thirds) do not display stable risk preferences across similar decision contexts.[1] Yet, Einav, Finkelstein, Pascu, and Cullen (2012) emphasize that there nevertheless is substantial commonality (rank correlation) in households' risky choices across contexts. This opens the question of whether there are other models, besides EU, that satisfy the stability criterion and can explain households' risky choices across different contexts.

Barseghyan, Molinari, O'Donoghue, and Teitelbaum (forthcoming) find that probability distortions — characterized by substantial overweighting of small probabilities and only mild insensitivity to probability changes — are a key feature of households' risk preferences.[2] Building on this finding, we study a model in which households behave roughly as in EU, except that event probabilities are systematically distorted. Such distortions may arise, for example, because of probability weighting, loss aversion, or disappointment aversion; see Barseghyan, Molinari, O'Donoghue, and Teitelbaum (forthcoming) for a discussion.

Using data on households' deductible choices in auto collision, auto comprehensive, and homeowners insurance, we first revisit the stability question and find that a model with linear probability distortions yields preference stability across all three contexts for the vast majority of households whose choices can be rationalized in a single context. Moreover, we find that the level of stability achieved by the probability distortion model is substantially higher than the level of stability that can be achieved by any of the following alternative models: an EU model with a variety of vNM utility functions, a probability weighting model with either Kahneman-Tversky or Prelec weighting functions, and a model with either Kőszegi-Rabin loss aversion or Gul disappointment aversion.

We then address the question of how much we can infer about the underlying structure of risk preferences by exploiting the stability criterion. We take a partial identification approach to answer this question; see, e.g. Manski (2003). Specifically, we make minimal assumptions about the structure of preferences and add these assumptions sequentially in order to transparently show the role that each plays in sharpening the inference. We make no distributional assumptions about the underlying unobserved heterogeneity in risk preferences, nor do we postulate any functional depen-

[1] Barseghyan, Prince, and Teitelbaum (2011) study households' deductible choices in auto collision insurance, auto comprehensive insurance, and homeowners insurance. Einav, Finkelstein, Pascu, and Cullen (2012) study employees' workplace benefit choices over health insurance, drug insurance, dental insurance, short-term disability insurance, long-term disability insurance, and 401(k) investments.

[2] Barseghyan, Molinari, O'Donoghue, and Teitelbaum (forthcoming) also study households' deductible choices in auto collision, auto comprehensive, and homeowners insurance.

dence of the preference parameters on observables. The idea is simply to use revealed preference arguments to bound the preference parameters.[3] The more choices per household we observe, the more precise the inference we can draw about the household's preferences.

While the revealed preference arguments coupled with the stability criterion are powerful, within a probability distortion model they need to be accompanied by shape restrictions on the probability distortion function to be informative. This is because households make choices in different contexts, and in each context they face a different "objective" probability of experiencing a claim. Hence, in a given context revealed preference arguments yield a bound on the probability distortion function evaluated at a *specific* claim probability. Some structure on the probability distortion function, therefore, is needed to give identifying power to the stability assumption. We introduce such structure sequentially. Our first assumption is monotonicity — we merely require that the probability distortion function returns higher values for higher claim probabilities. This assumption cannot be rejected for more than 80% of "rationalizable" households. We then add the assumption that the probability distortion function is concave. This assumption cannot be rejected for more than 75% of rationalizable households and for the vast majority of households for whom the monotonicity assumption cannot be rejected. It turns out that convexity of the probability distortion function also cannot be rejected for a substantial fraction of households. The last two assumptions together imply linearity — i.e., the probability distortion function is given by $\Omega(\mu) = a + b\mu$, where μ is the objective claim probability and a and b are the intercept and slope parameters. The linearity assumption cannot be rejected for about 75% of rationalizable households.

We conclude our analysis by addressing two questions: (1) What single probability distortion function comes "closest" to yielding stability for all rationalizable households? and (2) What is the fraction of households whose choices can be rationalized by this single probability distortion function? To answer these questions we propose a point estimator that minimizes the average Euclidean distance to each household's set of "stable" probability distortion functions. We establish consistency and asymptotic normality of this estimator and apply it to our data. We find that (1) our "minimum distance" probability distortion function is remarkably similar to the "maximum likelihood" probability distortion function estimated in Barseghyan, Molinari, O'Donoghue, and Teitelbaum (forthcoming) (it exhibits substantial overweighting of small probabilities and only mild insensitivity to probability changes) and (2) the choices of about 20% of households can be fully rationalized by this single probability distortion function.

References

Barseghyan, L., F. Molinari, T. O'Donoghue, and J. C. Teitelbaum (forthcoming): "The Nature of Risk Preferences: Evidence from Insurance Choices," *American Economic Review*.

Barseghyan, L., J. Prince, and J. C. Teitelbaum (2011): "Are Risk Preferences Stable Across Contexts? Evidence from Insurance Data," *American Economic Review*, 101(2), 591–631.

Einav, L., A. Finkelstein, I. Pascu, and M. R. Cullen (2012): "How General are Risk Preferences? Choice under Uncertainty in Different Domains," *American Economic Review*, 102(6), 2606–2638.

Manski, C. F. (2003): *Partial Identification of Probability Distributions*. Springer Verlag, New York.

[3] Recall that in our setting, as in similar settings with discrete choices, observing a single choice generally does not allow us to point identify a household's preference parameter — it only narrows the parameter space to a particular region in which every point is consistent with the household's choice.