Submitted to *Econometrica*

COMMENT ON ‘ECONOMICS AND MEASUREMENT’

JESSE M. SHAPIRO
Harvard University and NBER

This stimulating presidential address makes a forceful case for using economic theory to develop new measurement tools, and for using the resulting data to further develop economic theory. I find little to disagree with in the address and so will focus my remarks on making a case for the complementary study of how economic agents perceive the economic environment.

As a reminder, the paper defines a measurement system via the relation

$$\mathbf{m} = g(\theta, \varepsilon)$$

where $\mathbf{m}$ is an observed measurement, θ is an unobserved latent object of interest, ε is an unobserved measurement error, and $g(\cdot, \cdot)$ is a function. A researcher wishes to use information about the unobserved latent object of interest θ to learn about an economic parameter ϕ. The relation between θ and ϕ is encoded in an economic system via the relation

$$F(\theta; \phi) = 0$$

where $F(\cdot; \cdot)$ is some function prescribed by economic theory.

As a concrete example modeled loosely on Ameriks et al. (2020), suppose that the researcher is interested in the coefficient of relative risk aversion ϕ of a single household whose preferences over final wealth take the constant relative risk aversion form. The researcher may be interested in ϕ for several reasons, including to evaluate the welfare effects of different policies that affect the distribution of final wealth.

Table I summarizes several approaches to learning about ϕ. The first two rows of the table consider what might be thought of as a “traditional” approach. Here there is a measure $\mathbf{m}$ of the household’s asset allocations, for example the amounts held in risky and safe

Jesse M. Shapiro: jesse_shapiro@fas.harvard.edu

assets. Standard results (e.g., Campbell and Viceira 2002, chapter 2, equation 2.25) imply a tight relationship between the optimal share of wealth θ devoted to risky assets and the coefficient of relative risk aversion ϕ.

Given knowledge of the perceived risk premium and the perceived volatility of the return on the risky asset, it is possible to invert the relationship to identify the coefficient ϕ from θ. This approach naturally requires the economic assumption that the household's asset allocation is optimal at the time of measurement, precluding for example that the household has failed to optimize due to inertia or rebalancing costs. This approach also requires some method of recovering the perceived risk premium and the perceived volatility of the return on the risky asset. For example, the researcher may suppose that the household has the same beliefs as would a sophisticated econometrician, enabling the researcher to use econometric estimates of the risk premium and volatility as a proxy for the household's beliefs.

Using the traditional approach thus involves contestable assumptions on the economic system. By contrast, the traditional approach justifies fairly tight structure on the measurement system, because the latent variable θ that we seek to learn has an objective value. If the researcher has administrative data on asset allocations (as in, e.g., Fagereng et al. 2020), then the researcher might reasonably suppose that θ is measured perfectly. If the researcher has survey data on asset allocations, then the researcher will likely need to invoke the presence of response error ε, but the researcher's assumptions about ε, and about the function $g(\cdot,\cdot)$, can be informed by observation, for example by matching survey and administrative data for a subset of survey respondents (e.g., Ferber et al. 1969).

The next row of the table entails a less traditional approach. Here, the researcher presents the household with a hypothetical scenario, for example specifying the risk premium and volatility, and asks the household to declare the household's preferred asset allocation $\mathbf{m}$. This approach may justify fairly tight structure on the economic system. Because the question is hypothetical, frictions such as inertia or rebalancing costs do not prevent a household from reporting their optimal allocation. And because the researcher can directly specify the risk premium and volatility of the return on the risky asset, it may be reasonable to suppose that these correspond to the household's beliefs about the hypothetical scenario.

At the same time, relying on hypothetical choices may entail more contestable assumptions about the measurement system. A priori we may be unsure about how the household interprets the survey question, how the household reasons about its answer, and how the answer relates to the way the household would behave were the hypothetical scenario to realize in reality. Because the reported asset allocation $\mathbf{m}$ concerns a hypothetical scenario, it is not possible to directly verify its relationship to the true fraction θ that a household would allocate to the risky asset were the hypothetical scenario to realize.

Of course, some means of directly interrogating the structure of the measurement system are still available. The researcher can ask the household about hypothetical scenarios that correspond closely to some real situations in which choices can be observed, or to situations with a clear best (e.g., dominant) choice. The researcher can then study which modes of questioning result in the closest correspondence between the household's response and the observed, or best, choice. Alternatively, the researcher can modify the elicitation, for example by introducing some decisions with real stakes, to shorten the distance between the response $\mathbf{m}$ and the behavior of interest θ.

These directions each have their limits, because the advantages of hypothetical choice data lie precisely in the researcher's ability to ask about many potentially high-stakes scenarios without actually instantiating them. But these examples suggest that it may be possible to learn about $g(\cdot,\cdot)$ and ε in a principled way.

The final row in the table entails an approach that I would consider highly nontraditional, namely asking the household to state its coefficient of relative risk aversion. Here the measure $\mathbf{m}$ is the stated coefficient of relative risk aversion and the latent variable θ is the true coefficient of relative risk aversion. The economic system is then trivial, $F(\theta;\phi) = \theta - \phi$. As Ameriks et al. (2020) write, “it would be ideal if survey respondents could accurately and directly report their preference parameters” (p. 2395). But, it is very difficult to know what structure to put on the measurement system. Ameriks et al. (2020) continue, “of course, we cannot ask survey respondents to report their coefficient of relative risk aversion...” (p. 2395).

Comparing across the rows of Table I thus illustrates a tradeoff: taking advantage of less traditional forms of measurement may allow researchers to make fewer contestable economic assumptions but may entail making more contestable measurement assumptions.

One way to make the tradeoff shallower is to develop a better understanding of how economic agents reason about, express, and interpret the latent objects θ that play a direct role in economic models. There are promising developments underway in these areas, with economists documenting systematically the role of complexity in decision-making (e.g., Enke and Graeber forthcoming; Oprea 2022), the way individuals reason about the economy (e.g., Andre et al. 2022), and the way that different descriptions of an economic environment can influence behavior in that environment (e.g., Gonczarowski et al. 2022). Importantly, though in different contexts, all of these projects focus not (only) on how agents behave in different environments, but how agents perceive those environments, and how agents' perceptions shape their behavior. An understanding of such matters seems essential for taking best advantage of the kinds of measurement tools discussed in the address.

The value of such an understanding is suggested in the quotes from Ameriks et al. (2020). From the researcher's standpoint, the most valuable information to elicit from the agent is the economic parameter ϕ. But eliciting ϕ directly may entail representing the environment in a way that is very different from the agent's own representation. A traditional approach to closing this gap—exemplified by Friedman and Savage's (1948) billiards player—is to focus on the implications of the parameter for observable economic decisions under the researcher's model. A less traditional, complementary, and promising approach to closing the gap is for the researcher to learn more about the way the agent represents the environment.

REFERENCES

AMERIKS, JOHN, JOSEPH BRIGGS, ANDREW CAPLIN, MATTHEW D. SHAPIRO, AND CHRISTOPHER TONETTI (2020): “Long-Term-Care Utility and Late-in-Life Saving,” *Journal of Political Economy*, 128 (6), 2375–2451. [1, 3, 4]

ANDRE, PETER, INGAR HAALAND, CHRISTOPHER ROTH, AND JOHANNES WOHLFART (2022): “Narratives about the Macroeconomy,” *CEBI Working Paper 18/21*. [4]

CAMPBELL, JOHN Y. AND LUIS M. VICEIRA (2002): *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press. [2]

ENKE, BENJAMIN AND THOMAS GRAEBER (forthcoming): “Cognitive Uncertainty,” *Quarterly Journal of Economics*. [4]

FAGERENG, ANDREAS, LUIGI GUISO, DAVIDE MALACRINO, AND LUIGI PISTAFERRI (2020): “Heterogeneity and Persistence in Returns to Wealth,” *Econometrica*, 88 (1), 115–170. [2]

FERBER, ROBERT, JOHN FORSYTHE, HAROLD W. GUTHRIE, AND E. SCOTT MAYNES (1969): “Validation of Consumer Financial Characteristics: Common Stock,” *Journal of the American Statistical Association*, 64 (326), 415–432. [2]

FRIEDMAN, MILTON AND L. J. SAVAGE (1948): “The Utility Analysis of Choices Involving Risk,” *Journal of Political Economy*, 56 (4), 279–304. [4]

GONCZAROWSKI, YANNAI A., ORI HEFFETZ, AND CLAYTON THOMAS (2022): “Strategyproofness-Exposing Mechanism Descriptions,” *arXiv econ.TH*, 2209.13148. [4]

OPREA, RYAN (2022): “Simplicity Equivalents,” *UC Santa Barbara Working Paper*. [4]

TABLE I

APPROACHES TO LEARNING A HOUSEHOLD'S COEFFICIENT OF RELATIVE RISK AVERSION ϕ

Measurement m	Latent variable θ	Assumptions for measurement system $g(\cdot,\cdot)$	Assumptions for economic system $F(\cdot,\cdot)$
Administrative data on asset ownership.	Actual share of risky assets.	Perfect measurement.	Rational expectations, frictionless optimization
Self-reported data on asset ownership.	Actual share of risky assets.	Verifiable assumptions about reporting error.	Rational expectations, frictionless optimization.
Hypothetical choice data on asset purchases.	Hypothetical share of risky assets.	Partly verifiable assumptions about reporting error.	Optimization.
Direct elicitation of coefficient of relative risk aversion.	Coefficient of relative risk aversion.	Unverifiable assumptions about reporting error.	Minimal.