Decision Theory, Subjective Uncertainty, and Computer Science

GIACOMO LANZANI
Harvard University

I argue that further integration between Decision Theory and the methods of quantifying complexity and evaluating performance in Computer Science is valuable. I review [Lanzani 2024] as an illustration of this combination.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

General Terms: Economics, Performance, Theory

Additional Key Words and Phrases: Learning, Misspecification, Ambiguity, Convergence

1. OBJECTIVELY NORMATIVE DECISION THEORY

Decision Theory, as a subliterature within Economics, has always faced a tension between a normative interpretation and a descriptive one, between whether the decision criteria and axioms proposed should be interpreted as desiderata or postulates about the actual behavior of economic agents. One area in which the normative interpretation has been prevalent is decision theory under uncertainty.

However, for Subjective Expected Utility (SEU) and Bayesianism as postulated by [Savage 1954] (see [Cerreia-Vioglio et al. 2013] for its "belief over models" interpretation pursued here), the normative justification is an internal, or subjective, one. Seeing the axioms corresponding to the decision criterion, the decision maker feels compelled to adhere to them. Still, they do not imply that a decision-maker following them will perform particularly well in an objective environment. In static decision problems, there is no reason to expect such a good objective performance, as the decision maker is not omniscient. However, in dynamic environments, the gap between subjectively and objectively good decisions is elegantly closed for Bayesianism by Bernstein von Mises type theorems (cf. [Doob 1949; Breiman et al. 1964]), which show that if a Bayesian statistician is correctly specified, and the set of models considered is well-behaved, they will eventually concentrate their beliefs on the true data-generating process (DGP). Consequently, a correctly specified Bayesian agent will eventually take the ex-post optimal course of action.

In subsequent years, several different decision criteria have been proposed and axiomatized for decision theory under uncertainty. Many of these decision criteria have been motivated with an internally normative motivation; see the material reviewed in [Gilboa 2009; Gilboa and Marinacci 2016; Hansen and Marinacci 2016]. Chief among the normative motivations for the departure from subjective expected utility was a concern for complexity: many decision problems faced by economic agents are so complicated that it is impossible to quantify them with a set of proba-

Author's address: giacomolanzani34@gmail.com. I thank Roberto Corrao, Federico Echenique, Drew Fudenberg, and Irene Lo for their helpful comments.

bilistic models that could be fully trusted. Concretely, this impossibility comes from the fact that a set of relatively simple (e.g., parametric) models will be very prone to misspecification, i.e., not to include the true DGP. In contrast, an excessively large set of models will make Bayesian updating inconsistent (cf. [Diaconis and Freedman 1986]), or less dramatically, it will make convergence excessively slow.

However, so far, less progress has been made on whether these decision criteria *objectively* guard against the problems that have motivated their introduction. This question can be asked for countless decision criteria under uncertainty, e.g., the maxmin model [Gilboa and Schmeidler 1989], Choquet Expected Utility [Schmeidler 1989], multiplier preferences [Hansen and Sargent 2001; Strzalecki 2011], variational preferences [Maccheroni et al. 2006], models that combine worst and best case scenarios [Ghirardato et al. 2004; Gul and Pesendorfer 2014] or even general uncertainty averse preferences [Cerreia-Vioglio et al. 2011].[1] Do these criteria perform well in an objective situation in which a single probability measure would have been misspecified? How do we even operationalize the request for good performance? Does the reduction to a smaller set of DGPs paired with caution in using them overcome the inconsistency (or slow consistency) problems faced by the Bayesian paradigm with an extremely large set of models? I next turn to the discussion of why the approaches developed in Computer Science could be useful in formalizing and answering these questions.

2. WHY COMPUTER SCIENCE CAN HELP

The Computer Science literature has a long tradition in three areas suited to interact well with the goal of providing an objective assessment of the performance of decision rules under uncertainty: i) Criteria for the *objective* evaluation of decisions in repeated decision problems; ii) An emphasis on the *speed* of convergence; iii) An emphasis on the *complexity* of algorithms and decision rules.

Objective Evaluation Criteria. First, the Computer Science literature can contribute by providing a variety of criteria used to evaluate the performance of decision rules. The most studied is the concept of no-regret dynamics (see, e.g., [Noam Nisan 2011]). This requirement proved particularly useful for game-theoretic settings. Still, the dual representation as a game against Nature featured by uncertainty-averse decision criteria, [Cerreia-Vioglio et al. 2011] suggests that such techniques could prove useful even when a single decision maker is involved.

A less widely-used objective valuation criterion that suits the problems faced in Decision Theory is the advice-augmented dynamic performance. Loosely speaking,

[1] Of course, a dynamic version of these criteria has been formulated and studied, and the descriptive predictions for a decision-maker using such criteria and learning have been explored (see, [Marinacci 2002], [Epstein and Schneider 2007], and [Battigalli et al. 2019] as well as the extended references in [Hansen and Sargent 2011] and [Ilut and Schneider 2022] for results with particular attention to their macroeconomic and finance applications). The emphasis of this article is on decision theory under uncertainty, but I believe a similar perspective can also be useful in evaluating whether theories of choice under risk built on the promise to protect from uncertainty in the evaluation as [Cerreia-Vioglio et al. 2015] or correlation (see, e.g., [Lanzani 2022] for a unifying treatment) objectively deliver such protections.

it considers settings in which the decision-maker has access to the repeated predictions of an algorithm or a model and requires that, as long as the algorithm is not corrupted or the model is correct, the DM obtains the optimal performance if they completely trusted the advice. This approach has been successfully applied in many settings (see, e.g., [Mahdian et al. 2012; Banerjee et al. 2022]). This evaluation criterion combines particularly well with the economic approach in that economic agents are often endowed, even under non-SEU decision criteria, with a set of (possibly simplified) models they use to inform decisions.

Speed of Convergence. Second, the Computer Science literature pays special attention to the quantitative determination of the performance of a dynamic strategy. No regret, especially in individual problems, is hardly a goal for computer science but rather a given. The emphasis is more on determining the speed at which the regret decays to 0 (see, e.g., [Cesa-Bianchi and Lugosi 2006]).

Again, this focus seems to fit well with Decision Theory under uncertainty, especially given its motivation; see, e.g., this passage from the survey in [Gilboa 2023]:

"... my personal view is that it is wrong to assume that the only rational way to make decisions is to adopt a prior and follow SEUT ... If there is an infinite horizon of learning periods ahead of us, the choice of such a prior may be almost immaterial: as long as the prior is sufficiently open-minded, the underlying process would be learnt. But there are too many problems, ranging from wars to climate change, where we simply don't have the time to learn the underlying process ... In these cases, I believe that it may be more rational to admit that we do not know the probabilities than to pretend that we do."

But then, if non-SEU decision criteria are proposed as responses to excessively slow learning, their dynamically accumulated loss should be quantified and compared with the speed of learning with a large SEU model, a task for which the Computer Science tools seem particularly apt. These trade-offs are also at the core of the machine learning literature that proposed notions of the complexity of models (like Vapnik–Chervonenkis dimension and Rademacher complexity) that connect the speed of convergence to model complexity.[2]

Complexity. Third, uncertainty-averse decision criteria are often defended with a complexity argument. The idea is that postulating a comprehensive, all-encompassing model of the economic environment and performing dynamic optimization and learning within such a large model is computationally infeasible. Therefore, it is informally claimed we should use simpler scenario(s) but cautiously use them. However, these cautious criteria are often highly nonlinear, and thus, the corresponding optimization processes become more involved.

Although many papers illustrate the tractability of these problems in specific settings or under particular parametric assumptions, almost no work deals explicitly with the tradeoff between the complexity of Bayesian updating in a large SEU-Bayesian model versus the complexity in the optimization problems that depart

[2]The study of these tradeoffs dates at least back to [Valiant 1984], although most of the emphasis has been on avoiding the problem of overfitting.

from SEU. Here, computer science seems again to be able to provide relevant insights, as the quantification of complexity has always been a central topic. Indeed, recent years have seen contributions in this direction, see [Echenique et al. 2011; Fudenberg et al. 2022; Camara 2022].

3. ILLUSTRATION: [LANZANI 2024]

This section briefly summarizes the findings of [Lanzani 2024]. There, I take a class of *static* decision criteria under uncertainty motivated by robustness to misspecification [Hansen and Sargent 2001; Cerreia-Vioglio et al. 2022]. I then combine it with some objective evaluation criteria borrowed from the computer science literature, specifically the requirement of a minmax guarantee and an adaptation of the "advice-augmented" consistency mentioned in Section 2. I show that this requirement pins down a unique *dynamic* way to adjust those decision criteria in the face of accumulated evidence. I leverage this normative result to characterize the long-run behavior we should expect when agents employ those dynamically adjusted decision rules and when they depart from it in the direction of being excessively demanding or lenient in evaluating their model performance. The subsequent summary borrows extensively from [Lanzani 2023].

I consider an agent that repeatedly chooses among actions $a \in A$ that induce an unknown distribution over outcomes $y \in Y$. An objective data-generating process maps current actions into a probability distribution over outcomes $(p^*_a)_{a\in A}$. The agent does not know p^* but rather envisions a simpler set of models Q, where each $q = (q_a)_{a\in A} \in Q$ is also an action-contingent collection of probability distributions over outcomes. The agent has a utility function u over the joint realization of actions and outcomes. At every period, this choice is determined by maximizing the average (with respect to a belief $\mu \in \Delta(Q)$) of robust control assessments, where each assessment uses a different structured model as the benchmark:

$$\int_Q \min_{p_a\in\Delta(Y)} \left(\mathbb{E}_{p_a} \left[u\left(a, y\right) \right] + \frac{1}{\lambda} R\left(p_a || q_a\right) \right) \mathrm{d}\mu\left(q\right). \tag{1}$$

Here, R denotes the relative entropy, λ is a real number, and the decision maker trades off between the performance of the action under the conjectured models in Q and worst-case scenarios that are not too far in terms of relative entropy.

I introduce endogeneity in the misspecification concern: the better the structured models explain the past, the less concerned the agent is. I first establish a normative result. If the agent wants, across a large set of environments, to be guaranteed to both:

(1) Always achieve the minmax payoff;
(2) Achieve the (approximately) ex-post optimal payoff if their model is (approximately) correct,

then, they should evaluate their model using a log-likelihood ratio and keep the level of concern for misspecification proportional to this log-likelihood ratio. I also show a partial converse, in that every rule of adjustment of the concern for misspecification that is either asymptotically faster or slower than this fails one of these two properties. Observe that requirement (2) can be read as an approximate version of the advice-augmented dynamic performance criterion, where predictive algorithms are replaced by predictive economic models.

I then move to the descriptive analysis, where I allow departures from this normative benchmark and consider agents that are too demanding in evaluating the models' performance. Similarly, I allow the opposite case in which the agent is too lenient in evaluating their model and attributes too much unexplained evidence to sampling variability.

With this, I characterize the possible long-run behavior of these different types of misspecified agents. The limit actions of the lenient type must converge to a Berk-Nash [Esponda and Pouzo 2016; **?**] equilibrium, i.e., to an SEU best reply to beliefs supported on the models closest to the data-generating process. In mathematical terms, the limit action a must satisfy the following fixed point condition:

$$a \in \operatorname*{argmax}_{a' \in A} \int_Q \mathbb{E}_{q_{a'}} \left[u\left(a', y\right)\right] \mathrm{d}\nu\left(q\right)$$

with

$$\operatorname{supp} \nu \subseteq \operatorname*{argmin}_{q \in Q} R(p_a^* || q_a).$$

Instead, overemphasis on the model's failures in explaining the data by the demanding type induces convergence to a maxmin best reply to the absolutely continuous models with respect to the true one. In formal terms

$$a \in \operatorname*{argmax}_{a' \in A} \min_{p : \exists q \gg p} \mathbb{E}_{p_{a'}} \left[u\left(a', y\right)\right].$$

In contrast, a statistically sophisticated type maintains a non-trivial concern for misspecification. If their behavior converges, it converges to a robust control best reply to the models closest to the actual data-generating process. Moreover, the misspecification concern is endogenously determined by how well the best models fit the evidence generated by the limit action. In formal terms,

$$a \in \operatorname*{argmax}_{a' \in A} \int_Q \min_{p_{a'} \in \Delta(Y)} \left(\mathbb{E}_{p_{a'}} \left[u\left(a', y\right)\right] + \frac{1}{\lambda} R\left(p_{a'} || q_{a'}\right) \right) \mathrm{d}\nu\left(q\right)$$

with

$$\operatorname{supp} \nu \subseteq \operatorname*{argmin}_{q \in Q} R(p_a^* || q_a)$$

and

$$\lambda \simeq \operatorname*{argmin}_{q \in Q} R(p_a^* || q_a).$$

Finally, I point out that an endogenous concern for misspecification induces cycles, showing how my model could be used to rationalize the cyclical monetary policy documented in [Sargent 2008].

4. FINAL REMARKS

The comments in Section 2 may lead to the question of whether the elaborate decision criteria proposed by the literature on Decision Theory under uncertainty are needed, given the theoretical advance in Computer Science. If we can achieve, say, no regret, at an optimal speed under some algorithm, why bother, at least in dynamic environments, with normative Decision Theory?

In my view, there are two main reasons why this perspective is fallacious, both related to the speed of convergence. The environment faced by the economic agents will often change before the positive results of the Computer Science literature have time to apply. Similarly, the agents are often not completely impatient. In all these situations, using models that allow the agent to extrapolate from the consequences of one course of action to another (i.e., that do not treat the environment as a generalized bandit problem) is extremely valuable. At the same time, the complexity of the environment implies those models or scenarios will probably be incomplete and calls for combining them with robustness or cautious evaluations.

REFERENCES

Banerjee, S., Gkatzelis, V., Gorokh, A., and Jin, B. 2022. Online nash social welfare maximization with predictions. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1–19.

Battigalli, P., Francetich, A., Lanzani, G., and Marinacci, M. 2019. Learning and self-confirming long-run biases. *Journal of Economic Theory 183*, 740–785.

Breiman, L., LeCam, L., and Schwartz, L. 1964. Consistent estimates and zero-one sets. *The Annals of Mathematical Statistics 35,* 1, 157–161.

Camara, M. K. 2022. Computationally tractable choice. In *EC*. 28.

Cerreia-Vioglio, S., Dillenberger, D., and Ortoleva, P. 2015. Cautious expected utility and the certainty effect. *Econometrica 83,* 2, 693–728.

Cerreia-Vioglio, S., Hansen, L. P., Maccheroni, F., and Marinacci, M. 2022. Making decisions under model misspecification.

Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., and Montrucchio, L. 2011. Uncertainty averse preferences. *Journal of Economic Theory 146,* 4, 1275–1330.

Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., and Montrucchio, L. 2013. Classical subjective expected utility. *Proceedings of the National Academy of Sciences 110,* 17, 6754–6759.

Cesa-Bianchi, N. and Lugosi, G. 2006. *Prediction, learning, and games.* Cambridge university press.

Diaconis, P. and Freedman, D. 1986. On the consistency of bayes estimates. *The Annals of Statistics 14,* 1, 1–26.

Doob, J. L. 1949. Application of the theory of martingales. *Le calcul des probabilites et ses applications*, 23–27.

Echenique, F., Golovin, D., and Wierman, A. 2011. A revealed preference approach to computational complexity in economics. In *Proceedings of the 12th ACM conference on Electronic commerce*. 101–110.

Epstein, L. G. and Schneider, M. 2007. Learning under ambiguity. *The Review of Economic Studies 74,* 4, 1275–1303.

Esponda, I. and Pouzo, D. 2016. Berk–nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica 84,* 3, 1093–1130.

Fudenberg, D., Kleinberg, J., Liang, A., and Mullainathan, S. 2022. Measuring the completeness of economic models. *Journal of Political Economy 130,* 4, 956–990.

Ghirardato, P., Maccheroni, F., and Marinacci, M. 2004. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory 118,* 2, 133–173.

Gilboa, I. 2009. *Theory of decision under uncertainty.* Number 45. Cambridge university press.

Gilboa, I. 2023. Decision under uncertainty state of the science. *Annual Review of Economics forthcoming.*

GILBOA, I. AND MARINACCI, M. 2016. Ambiguity and the bayesian paradigm. In *Readings in formal epistemology: Sourcebook*. Springer, 385–439.

GILBOA, I. AND SCHMEIDLER, D. 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18,* 2, 141–153.

GUL, F. AND PESENDORFER, W. 2014. Expected uncertain utility theory. *Econometrica 82,* 1, 1–39.

HANSEN, L. P. AND MARINACCI, M. 2016. Ambiguity aversion and model misspecification: An economic perspective. *Statistical Science 31,* 4, 511–515.

HANSEN, L. P. AND SARGENT, T. J. 2001. Robust control and model uncertainty. *American Economic Review 91,* 2, 60–66.

HANSEN, L. P. AND SARGENT, T. J. 2011. Robustness. In *Robustness*. Princeton university press.

ILUT, C. L. AND SCHNEIDER, M. 2022. Modeling uncertainty as ambiguity: A review.

LANZANI, G. 2022. Correlation made simple: Applications to salience and regret theory. *The Quarterly Journal of Economics 137,* 2, 959–987.

LANZANI, G. 2023. Dynamic concern for misspecification. In *Proceedings of the 24th ACM Conference on Economics and Computation*.

LANZANI, G. 2024. Dynamic concern for misspecification.

MACCHERONI, F., MARINACCI, M., AND RUSTICHINI, A. 2006. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica 74,* 6, 1447–1498.

MAHDIAN, M., NAZERZADEH, H., AND SABERI, A. 2012. Online optimization with uncertain information. *ACM Transactions on Algorithms (TALG) 8,* 1, 1–29.

MARINACCI, M. 2002. Learning from ambiguous urns. *Statistical Papers 43,* 1, 143–151.

NOAM NISAN, TIM ROUGHGARDEN, E. T. V. V. V. 2011. *Algorithmic Game Theory*. Cambridge University Press.

SARGENT, T. J. 2008. Evolution and intelligent design. *American Economic Review 98,* 1, 5–37.

SAVAGE, L. J. 1954. *The foundations of statistics*. Courier Corporation.

SCHMEIDLER, D. 1989. Subjective probability and expected utility without additivity. *Econometrica: Journal of the Econometric Society*, 571–587.

STRZALECKI, T. 2011. Axiomatic foundations of multiplier preferences. *Econometrica 79,* 1, 47–73.

VALIANT, L. G. 1984. A theory of the learnable. *Communications of the ACM 27,* 11, 1134–1142.