Games and Economic Behavior 100 (2016) 69–94

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Learning the Krepsian state: Exploration through consumption ☆

Evan Piermont[a], Norio Takeoka[b], Roee Teper[a,*]

[a] University of Pittsburgh, Department of Economics, United States
[b] Hitotsubashi University, Department of Economics, Japan

ARTICLE INFO

Article history:
Received 18 January 2016
Available online 21 September 2016

JEL classification:
D81
D83

Keywords:
Responsive learning
Subjective learning
Conditional preferences
Preference for flexibility

ABSTRACT

We take the Krepsian approach to provide a behavioral foundation for responsive subjective learning processes. In contrast to the standard subjective state space models, the resolution of uncertainty regarding the true state is endogenous and depends on the decision maker's actions. There need not be full resolution of uncertainty between periods. When the decision maker chooses what to consume, she also chooses the information structure to which she will be exposed. When she consumes outcomes, she learns her relative preference between them; after each consumption history, the decision maker's information structure is a refinement of the previous information structure. We provide the behavioral restrictions corresponding to a recursive representation exhibiting such a learning process. Through the incorporation of dynamics we are able to identify the set of preferences the decision maker believes possible after each history of consumption, without appealing to an environment with risk.

1. Introduction

Uncertainty regarding the state of the world, or one's preferences, underlies almost every economic environment. In such settings, the agent might have the opportunity to actively explore in order to increase her understanding. A typical example are the models of strategic experimentation (Robbins, 1952; Gittins and Jones, 1979; Weitzman, 1979).[1] When the agent takes an action, she does not only derive utility from it, but also observes its consequence and gains better information pertaining to the underlying uncertainty. It is intuitive, then, that the agent's understanding (or lack thereof) is a function of her experience, and subsequently, of her previous choices. We refer to such learning models as *responsive*. This paper provides a behavioral foundation and identification techniques for a model of *subjective and responsive learning*.

Kreps (1979) and Dekel et al. (2001) put forth the canonical models of introspective uncertainty, which identify the subjective uncertainty the decision maker (henceforth, DM) is facing regarding her own preferences. Implicitly, these models are static and assume that all uncertainty is realized in a way that is independent of the decision made by the agent. More recently, Takeoka (2007) and Dillenberger et al. (2014) introduce (static) models of gradual learning, and Krishna and Sadowski (2014, 2015) provide the behavioral foundation for a dynamic version of Kreps' model that allows for tastes

☆ The authors wish to thank Kevin Cooke, Eddie Dekel, Kfir Eliaz, Edi Karni, R. Vijay Krishna, Pietro Ortoleva, Wolfgang Pesendorfer, Philipp Sadowski, Marciano Siniscalchi, Francois Forges, an Associate Editor and two anonymous referees for their comments and remarks. Takeoka gratefully acknowledges financial support from a Grant-in-Aid for Scientific Research (C).

* Corresponding author.
E-mail addresses: ehp5@pitt.edu (E. Piermont), norio.takeoka@r.hit-u.ac.jp (N. Takeoka), rteper@pitt.edu (R. Teper).

[1] See Bergemann and Välimäki (2006) for a more recent survey.

http://dx.doi.org/10.1016/j.geb.2016.09.002
0899-8256/

shocks. However, in these models, as in Kreps, learning is not *responsive*: the arrival of information is irrespective of the consumption choices made by the decision maker.

In many circumstances, such as a DM acknowledging the possibility of transient taste shocks, this is perfectly reasonable. It is likely the DM's choice of restaurant in the morning will not affect her ability to discern her preference for fish and steak in the evening. However, if we consider the scenario of a DM learning her ranking over as-of-yet unconsumed alternatives, unresponsive learning is inadequate. For instance, imagine a novice researcher who is deciding on a project, but who is uncertain about her particular talents. She obtains no information without experience; there is no "exogenous" signal that will indicate her skill set. If she explores different research directions it will become apparent in which area she is more talented.

Weitzman (1979) introduced a model that studies the optimal strategy of a decision maker in this situation. The agent is engaged in a dynamic problem, trying to understand which project to invest in. Once she explores a project, she knows exactly what that project yields; this may also convey some information about other projects (in case the projects' outcomes are correlated). At each point the agent has to decide whether to keep exploring new projects or to commit to already explored and known projects.

We here study the decision theoretic aspects of responsive subjective learning. We follow the decision theoretic learning literature and consider a dynamic constrained choice environment that extends the Krepsian framework. We provide the axiomatic foundation of a recursive utility function exhibiting subjective responsive learning. Initially, just as in Kreps, there is uncertainty regarding preferences. Then, each period, the DM jointly chooses a consumption outcome and a constraint for the following period so as to maximize her (current period) consumption utility and continuation value. The DM takes into account that her choice of consumption today may teach her about her preferences, altering her information structure, and accordingly, her preferences over future constraints. This representation captures the notion of responsive learning by allowing the continuation utility function and information structure to explicitly depend on previous consumption choices. In this respect, our framework expresses the intertemporal tradeoff between the consumption value of the choice today and its future informational value.

The decision maker's subjective information structure changes *only* in response to her consumption choices, in a manner similar to Weitzman (1979).[2] In particular, our model considers a learning process characterized in part by two constraints. First, the DM's uncertainty regarding her ranking between a and b is resolved once a and b are consumed. Second, if the DM is initially uncertain about her ranking between a and b, then this uncertainty is not fully resolved unless she consumes outcomes over which her preferences are perfectly correlated with a and b.

1.1. Results

The learning characteristics described above imply that preferences over already consumed goods are stationary. We begin the analysis by considering a somewhat simple framework combining flexibility and dynamics, where the DM chooses between menus of streams of outcomes. In this framework we state Theorem 1, showing that stationarity in the presence of flexibility is sufficient in order to uniquely identify the subjective uncertainty. It is worth noting, uniqueness is achieved even though we do not consider environments with risk, as is customary in the literature beginning with the uniqueness result provided by Dekel et al. (2001).

We then proceed to develop a complete model of learning in a framework of dynamic programing. In this domain we state and prove Theorem 3, which is the axiomatization of responsive subjective learning. In particular, we are able to elicit from behavior, not only (1) the set of states the DM believes possible, but also (2) how the DM expects to learn conditional on a given path of consumption, and (3) how the decision maker anticipates her preferences, over both future consumption and consequent information, will change after learning any one of the pieces of information she believes possible. The elicitation of (2) and (3) are novel.[3] The main complication in accommodating responsive learning is identifying (3).

What makes this elicitation, and thus the axiomatization, inherently difficult, is that we are interested in states that are not only subjective but also conditional. To understand the DM's conditional preferences for information, it is first necessary to be able to condition on the relevant (subjective) state. This is not straightforward; the information structure is not directly observable as it is not incorporated into the primitive. A contribution of this paper is developing the tools allowing the modeler to elicit preferences conditional on subjective states.

Finally, Theorem 4 states that the subjective information structure (that is, the initial underlying uncertainty and how the DM expects to learn as she consumes) in our model is uniquely identified. What facilitates the identification in our environment is the dynamic and cumulative structure of learning. Our framework elicits the DM's (anticipated) understanding of uncertainty at multiple points in the learning process. It is the stationarity of preferences over already consumed goods,

[2] In a more general bandit problem, when learning pertains to distributional parameters, our partitional learning seems perhaps less realistic and noisy signals would be a natural generalization. It is also interesting to think about models that combine responsive learning and exogenous information flows (the latter as in Dillenberger et al., 2014). These models bear important economic content, but their axiomatizations are far from being close to the one suggested here. For example, Piermont and Teper (2016) axiomatize model of general responsive learning, albeit in a substantially different choice environment.

[3] The concept of conditional preferences of conditional preferences, etc., has been examined in other contexts. For examples see, Gul and Pesendorfer (2005) and Siniscalchi (2011).