Utility Maximization Under Endogenous Uncertainty

Ayush Gupta*

Boston University

Last Updated: 10 June 2025

Abstract

This paper establishes a general existence result for expected utility maximization in settings where the agent's decision affects the uncertainty faced by her. We introduce a continuity condition for choice-dependent probability measures which ensures the upper semi-continuity of expected utility. Our topological proof imposes minimal restrictions on the utility function and the random variable. In particular, we do not need common assumptions like the monotone likelihood ratio property (MLRP) or the convexity of distribution functions condition (CDFC). Additionally, we identify sufficient conditions – continuity of densities and stochastic dominance – which help verify our assumptions in most practical applications. These findings expand the applicability of expected utility theory in settings with endogenous uncertainty.

*I am indebted to Alex Chan and Mark Kon for their help and support throughout this project.
Email: guptaayu@bu.edu

1 Introduction

The existence of an expected utility maximizing action is a minimal requirement for any model of decision-making under uncertainty. Kennan (1981) showed that such an optimal action exists under mild continuity and compactness conditions. However, Kennan's result assumes that the uncertainty in the environment is captured by a random variable with a fixed distribution. Many economic settings violate that assumption: lifestyle choices affect an individual's health outcomes, R&D spending affects a firm's profits, and the level of effort affects a worker's wage. In each of these examples, the agent's action (lifestyle / investment / effort) alters the distribution of a utility-relevant random variable. The importance of such endogenous uncertainty is highlighted, for example, by Hansen and Sargent (2001). They explicitly consider environments where agents' decisions affect the underlying uncertainty but simply assume the existence of an optimal policy. Similar implicit assumptions are widespread across economics. The goal of this paper is to fill that gap by extending existence theorems to explicitly accommodate decision-dependent uncertainty.

The analysis presented here is most closely related to certain existence results in the principal–agent literature. In many of those models, the contract induces a mapping from the agent's choice of effort to a distribution over wages. Early work by Holmström (1979) and Shavell (1979) established the existence of optimal actions under mild assumptions. However, subsequent research, notably by Rogerson (1985), showed that much stronger regularity assumptions are typically needed. In particular, the widely-used first-order approach requires the monotone likelihood ratio property (MLRP) and the convexity of distribution functions condition (CDFC). These assumptions ensure that an optimal action exists but are restrictive and often violated in more general settings.

This paper does not impose either assumption. Instead, we introduce an upper semi-continuity condition for choice-dependent probability measures. A topological proof is then used to show that expected utility is upper semi-continuous and an optimal action exists. We also establish sufficient conditions that are easier to verify and ensure that our main assumption is satisfied. Ultimately, our approach yields a more general existence result and allows expected utility theory to be applied in a wider class of models with decision-dependent uncertainty.

2 Notation and Assumptions

Let A be the set of alternatives available to the agent. Let Ω be the (possibly infinite) set of possible realizations of a random variable. The associated probabilities are given by the measure m_a. The subscript $a \in A$ signifies that the measure depends on the agent's choice. Note that we do not require the support of the random variable to be the same for every choice of a. Concretely, we allow for $\omega \in \Omega$ such that $m_a(\omega) = 0$ but $m_{a'}(\omega) > 0$.

Let $u(a, \omega)$ be the agent's utility function. We want to prove that the following maximization problem has at least one solution:

$$\max_a v(a) = \max_a \mathbb{E}\big[u(a, \omega)\big]$$

This problem can be written as:

$$\max_a v(a) = \max_a \int_\Omega u(a, \omega)\, dm_a(\omega)$$

Our result requires the following assumptions:

1. A is a compact and first countable topological space.

2. $u(\cdot, \omega)$ is a set of equicontinuous and real valued functions on A.

3. $u(a, \cdot)$ is a real valued random variable for all $a \in A$ and is integrable.

4. For any function $g(\omega)$ which is integrable with respect to m_a and any sequence $a_n \to a$ the measure satisfies:

$$\int_\Omega g(\omega)\, dm_a(\omega) \geq \limsup_n \int_\Omega g(\omega)\, dm_{a_n}(\omega)$$

Assumptions 1 and 3 are ubiquitous in the literature. Assumption 2 is a slightly stronger version of the standard continuity assumption. While it may appear restrictive at first, we point out that Assumption 2 will be satisfied if, for example, Ω is compact and $u(a, \omega)$ is continuous in ω. These two assumptions are very common in the literature and are satisfied by many economic problems.

Assumption 4 is our main assumption. It imposes the notion of upper semi-continuity on the mapping from A to the space of probability measures. We cannot use standard definitions here because each probability measure is itself a mapping from the subsets of Ω to the real line. Moreover, it is not enough for the probability of individual events to vary continuously. Because expected utility depends on the entire distribution, the measure as a whole needs to change in a well-behaved manner.[1]

[1] Assumption 4 is much weaker than the standard $\epsilon - \delta$ definition of upper semi-continuity under the sup norm.

This paper does not prove necessity but it is easy to see that the existence of a utility maximizing action is not guaranteed without some version of Assumption 4. We use examples to illustrate the importance of this assumption and argue that it fails precisely when an optimal action need not exist. We also establish sufficient conditions which are easier to verify and ensure that Assumption 4 holds.

Note that our assumptions do not impose monotonicity or concavity on $v(a)$. Moreover, we do not need to assume that $u(a, \omega)$ is bounded on $A \times \Omega$.[2] This makes our result broadly applicable and ensures existence in many settings which are not covered by existing results.

3 Results

Our main result can be stated as:

Theorem 1. The expected utility maximization problem has at least one solution if Assumptions 1 - 4 are satisfied.

Proof. We have assumed that A is compact (Assumption 1). So a solution to the maximization problem is guaranteed if we can show that $v(a)$ is upper semi-continuous. That requires the following for any $a_n \to a$:

$$\int_\Omega u(a, \omega)\, dm_a(\omega) \geq \limsup_n \int_\Omega u(a_n, \omega)\, dm_{a_n}(\omega)^\dagger$$

The right hand side of the inequality can be written as:

$$\begin{aligned} \limsup_n \int_\Omega u(a_n, \omega)\, dm_{a_n}(\omega) \leq & \limsup_n \int_\Omega u(a, \omega)\, dm_{a_n}(\omega) \\ & + \limsup_n \int_\Omega \Big[u(a_n, \omega) - u(a, \omega)\Big]\, dm_{a_n}(\omega) \end{aligned} \tag{1}$$

Applying Assumption 4 to the first term in (1) gives:

[2]Unbounded utility is not incompatible with equicontinuity. Both are possible if, for example, the utility function is additively separable.

†Notice that the integrand and the measure are both changing with n. This dual dependence prevents us from applying standard results and necessitates the present approach.

$$\limsup_n \int_\Omega u(a,\omega)\, dm_{a_n}(\omega) \leq \int_\Omega u(a,\omega)\, dm_a(\omega)$$

Recall that $u(\cdot,\omega)$ is an equicontinuous set of functions on a compact space (Assumptions 1 and 2). Therefore $u(\cdot,\omega)$ are uniformly equicontinuous. Then:

$$\epsilon_n := \sup_\omega \Big| u(a_n,\omega) - u(a,\omega) \Big| \to 0$$

Applying this to the second term in (1) gives:

$$\begin{aligned}
\limsup_n \int_\Omega \Big[u(a_n,\omega) - u(a,\omega) \Big]\, dm_{a_n}(\omega) &\leq \limsup_n \int_\Omega \sup_\omega \Big| u(a_n,\omega) - u(a,\omega) \Big|\, dm_{a_n}(\omega) \\
&= \limsup_n \int_\Omega \epsilon_n \, dm_{a_n}(\omega) \\
&= \limsup_n \epsilon_n m_{a_n}(\Omega) \\
&= 0
\end{aligned}$$

Combining everything gives:

$$\begin{aligned}
\limsup_n \int_\Omega u(a_n,\omega)\, dm_{a_n}(\omega) &\leq \limsup_n \int_\Omega u(a,\omega)\, dm_{a_n}(\omega) \\
&\quad + \limsup_n \int_\Omega \Big[u(a_n,\omega) - u(a,\omega) \Big]\, dm_{a_n}(\omega) \\
&\leq \int_\Omega u(a,\omega)\, dm_a(\omega) + 0
\end{aligned}$$

This proves that $v(a)$ is upper semi-continuous and completes our proof of existence.

$\square$

The importance of Assumption 4 is obvious from the proof of Theorem 1. However, it is often challenging to verify in practical applications. We therefore establish sufficient conditions which facilitate the verification of Assumption 4 in most economically relevant contexts.

Assume hereafter that Ω has a σ-finite measure μ such that m_a is absolutely continuous with respect to μ for every a. Then the Radon–Nikodym theorem tells us that each m_a induces a density function f_a. These densities (when they exist) can be used to show that Assumption 4 holds.

Proposition 1. Assumption 4 is satisfied if for every $a_n \to a$ and any $\epsilon > 0$ there exists N such that $\Big|f_a(\omega) - f_{a_n}(\omega)\Big| \leq \epsilon f_a(\omega)$ almost everywhere for $n > N$.

Proof. Let $g(\omega)$ be any function which is integrable with respect to $f_a(\omega)$.

Fix $\epsilon > 0$. Then $\Big|g(\omega)f_{a_n}(\omega)\Big|$ is dominated by $\Big|g(\omega)\Big|(1+\epsilon)f_a$ for sufficiently large n.

We can thus apply the dominated convergence theorem to get:

$$\begin{aligned}
\int_\Omega g(\omega)\, dm_a(\omega) &= \int_\Omega g(\omega) f_a(\omega)\, d\mu(\omega) \\
&= \int_\Omega \lim_n g(\omega) f_{a_n}(\omega)\, d\mu(\omega) \\
&= \lim_n \int_\Omega g(\omega) f_{a_n}(\omega)\, d\mu(\omega) \\
&= \limsup_n \int_\Omega g(\omega) f_{a_n}(\omega)\, d\mu(\omega) \\
&= \limsup_n \int_\Omega g(\omega)\, dm_{a_n}(\omega)
\end{aligned}$$

□

Remark 1. Proposition 1 only establishes a sufficient condition for Assumption 4. The given condition does not ensure the existence of a utility maximizing action. Nor does it make the proof of Theorem 1 trivial.

That is because $\Big|u(a_n, \omega) f_{a_n}(\omega)\Big|$ need not have a dominating function even when the condition in Proposition 1 is satisfied. Consequently, the dominated convergence theorem does not apply and we cannot move the limit inside the integral.

Remark 2. Proposition 1 requires that $f_a(\omega)$ satisfies a uniform continuity condition with respect to a – the choice variable. This is independent of continuity with respect to ω – the random variable. In particular, a family of discrete random variables may satisfy the condition in Proposition 1 while a family of continuous random variables does not.

For example, assume that ω is uniformly distributed on the interval $[0, a]$. Then the density $f_a(x)$ for any $x > 0$ is discontinuous at $a = x$.

Verifying the continuity of a family of density functions is typically much easier than verifying Assumption 4 directly. Proposition 1 thus allows us to prove that a utility maximizing action exists in many commonly used models.

The intuition from Proposition 1 also helps us identify situations in which an optimal action need not exist. Consider the following example:

Example 1.

$$a \leq 1 \Rightarrow \omega = \begin{cases} 0, & \text{with probability } \frac{1}{2} \\ 1, & \text{with probability } \frac{1}{2} \end{cases}$$

$$a > 1 \Rightarrow \omega = 1, \text{ with probability } 1$$

It is clear that the utility maximization problem need not have a solution if $u(a, \omega)$ is decreasing in a and increasing in ω. That is because the random variable "jumps" at $a = 1$. This jump causes *expected* utility to be discontinuous, even when the utility function is continuous in both arguments.

Such jumps in the random variable is precisely what is ruled out by Assumption 4.[3] However, intuition suggests that jumps are not always an impediment to expected utility maximization. We illustrate this with the following example:

Example 2.

$$a < 1 \Rightarrow \omega = \begin{cases} 0, & \text{with probability } \frac{1}{2} \\ 1, & \text{with probability } \frac{1}{2} \end{cases}$$

$$a \geq 1 \Rightarrow \omega = 1, \text{ with probability } 1$$

Example 2 is consistent with the decision problem faced by workers in many principal-agent models. If a is the worker's level of effort and ω is her wage, we expect at least one solution to exist. However, this random variable also jumps at $a = 1$ and it is easy to verify that Assumption 4 is not satisfied.

Nevertheless, this example does not undermine our result. On the contrary, it highlights a strength of Theorem 1. That is because, in this example, the existence of a solution is contingent on our choice of utility function. For instance, a solution may fail to exist if $u(a, \omega)$ is increasing in a and decreasing in ω. In contrast, Theorem 1 ensures the existence of a solution for any possible utility function.

Motivated by Example 2, we establish the sufficiency of a much weaker condition when additional assumptions are imposed on the utility function:

[3] Consider $g = \omega$. It is clear that $\int_\Omega g \, dm_a(\omega) \ngeq \limsup_n \int_\Omega g \, dm_{a_n}(\omega)$ for $a_n \to a = 1$ if $a_n > 1$ for every n.

Proposition 2. Assume $g(\omega)$ is monotone increasing in ω. Then Assumption 4 is satisfied if $\limsup_n f_{a_n}$ is (weakly) first order stochastically dominated by f_a for all $a_n \to a$.

Proof.

$$\begin{aligned}
\limsup_n \int_\Omega g(\omega) f_{a_n}(\omega) d\mu(\omega) &= \limsup_n \mathbb{E}\big[g(\omega_{a_n})\big] \\
&\leq \mathbb{E}\big[g(\omega_a)\big] \\
&= \int_\Omega g(\omega) f_a(\omega) d\mu(\omega)
\end{aligned}$$

The inequality comes from the fact that $g(\omega)$ is increasing in ω and ω_a first order stochastically dominates ω_{a_n}.

□

A symmetric argument shows that $\limsup_n f_{a_n}$ first order stochastically dominating f_a is sufficient when $g(\omega)$ is monotone decreasing in ω.

Proposition 2 significantly broadens the scope of our analysis. It can ensure the existence of a solution even when the random variable exhibits jumps, but at the cost of an additional assumption on the utility function. Applying this result to Example 2 confirms our intuition that a solution exists whenever $u(a, \omega)$ is increasing in ω.

4 Conclusion

Endogenous uncertainty is common in many economic settings. However, existing results require strong assumptions which are frequently violated in practice. The challenge is integrating over all possible states of the world when the integrand (utility) and the probability measure are both changing with the agent's action. This dual dependence precludes straightforward applications of Fatou's Lemma or conventional dominated convergence arguments. The main technical contribution of this paper is a topological proof demonstrating that, under the proposed continuity condition, expected utility remains upper semi-continuous. This work also establishes intuitive sufficient conditions – continuity of densities and stochastic dominance – which make our main assumption easier to verify in most applications. Through these contributions, this paper significantly advances the applicability of expected utility theory in models with endogenous uncertainty.

References

Lars Peter Hansen and Thomas J Sargent. Robust control and model uncertainty. *American Economic Review*, 91(2):60–66, 2001.

Bengt Holmström. Moral hazard and observability. *The Bell Journal of Economics*, 10(1):74–91, 1979.

John Kennan. The existence of expected utility maximizing decisions when utility is unbounded. *Econometrica*, 49(1):215–218, 1981.

William P Rogerson. The first-order approach to principal-agent problems. *Econometrica*, 53(6): 1357–1367, 1985.

Steven Shavell. Risk sharing and incentives in the principal and agent relationship. *The Bell Journal of Economics*, 10(1):55–73, 1979.