

Risk Aversion, Information Acquisition, and Technology Adoption

Canan Ulu

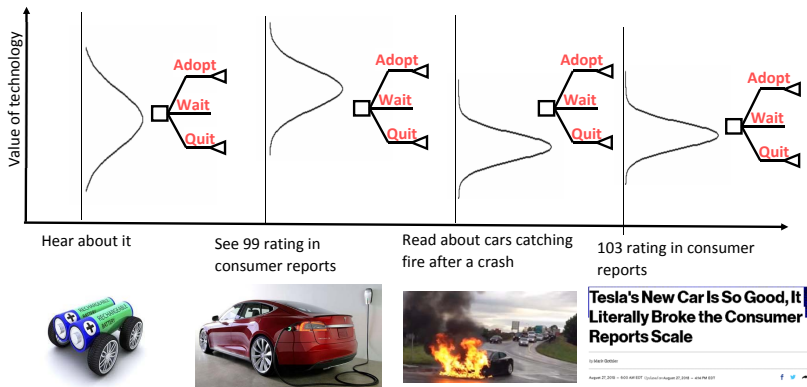
McDonough School of Business
Georgetown University

Jim Smith

Fuqua School of Business
Duke University

SAMSI Games and Decisions in Reliability and Risk Workshop
May 2016

Problem: Should Jim buy a Tesla? Or should he wait and learn more?



Other examples: farmer planting a new variety of soybean, utility building a power plant based on a new technology, or doctors changing treatments

We study a DP model of information acquisition in technology adoption decisions.

- We build on McCardle (1985) and Ulu and Smith (2009), adding *risk aversion*.
- In each period, the consumer can adopt the technology, gather information about the technology, or quit.

State variables: probability distribution on technology benefits
wealth

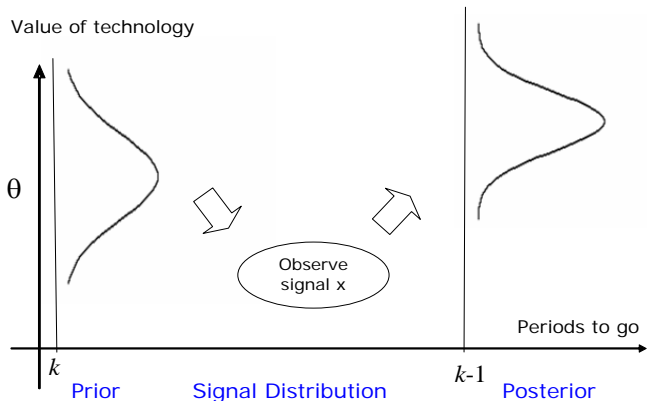
Beliefs are updated over time using Bayes' Rule.

Arbitrary distributions are allowed.

Information gathering is costly.

- We focus on structural properties of the model:
 - Properties of the value function (increasing, convex, ...)
 - Monotonicity properties of the optimal policies
 - Effects of risk aversion

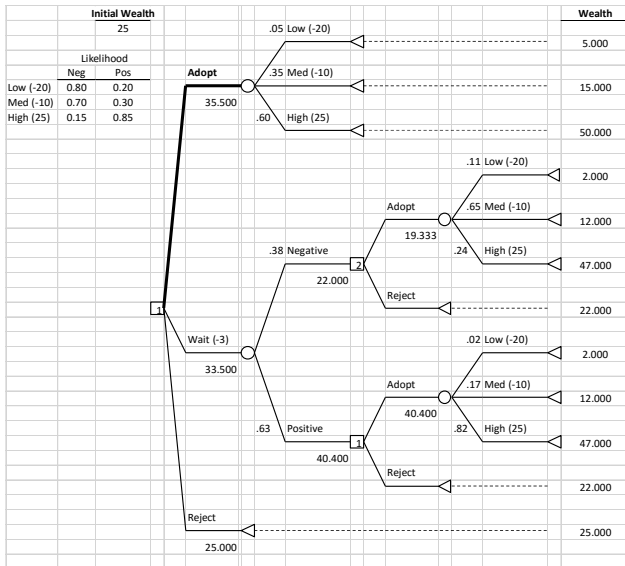
Modeling learning: Notation



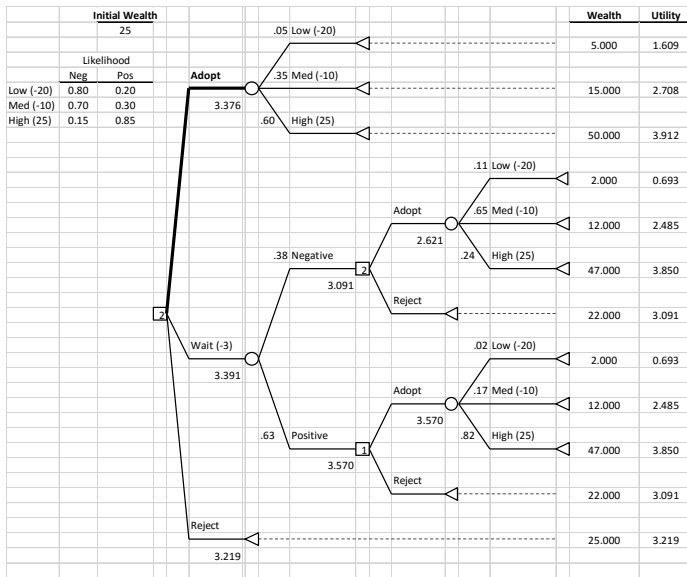
$$f(x; \theta) = \int_{\theta} L(x; \theta) \pi(\theta) d\theta$$

$$\pi(\theta; x) = \frac{L(x; \theta) \pi(\theta)}{\int L(x; \theta) \pi(\theta) d\theta}$$

Decision Tree Example: Risk Neutral



Decision Tree Example: Risk Averse



The value function:

- c = cost of waiting ($c > 0$)
- $u(w)$ = DM's utility for wealth w
- Value (or derived utility) function with k periods remaining:

$$U_0(w; \delta) = \max_{\text{adopt}} \left\{ u(w) \right\} \quad \text{(adopt)}$$

$$U_k(w; \delta) = \max_{\text{wait}} \left\{ E[U_{k-1}(w - c; (\delta; \kappa))] \right\} \quad \text{(wait)}$$

$$U_k(w; \delta) = \max_{\text{quit}} \left\{ u(w) \right\} \quad \text{(quit)}$$

where

$$E[u(w + \tilde{y})] = \int_{-\infty}^{\infty} u(w + y) f(y) dy$$

$$E[U_{k-1}(w; (\delta; \kappa))] = \int_{-\infty}^{\infty} U_{k-1}(w; (\delta; x)) f(x) dx$$

Illustrative example: beta-Bernoulli model

- $\theta = p \in [0, 1]$ where $f(p) = p^{\alpha-1}(1-p)^{\beta-1}$ with $p \in [0, 1]$
- Expected benefit $E[\theta] = \frac{\alpha}{\alpha + \beta}$; "precision" = $(\alpha + \beta)$
- Signals are \pm or $-$ with probability p or $(1-p)$. Precision increases by one each period. Given prior with $(\alpha; \beta)$,
 - + signal $\Rightarrow (\alpha + 1; \beta)$
 - signal $\Rightarrow (\alpha; \beta + 1)$
- Example: Start with $(\alpha; \beta) = (2; 25)$, observe $(+; +; -; -)$:

Illustrative example: risk-neutral results

- $u(w) = w$; initial wealth = 1.06; $c = 0.01$; long time horizon

Policy regions

Value function with $\gamma = 10$

Illustrative example: risk-averse results

- $u(w) = 1 - 0.2w^{(1-\gamma)}$ where $\gamma = 6$; initial wealth = 1.06; $c = 0.01$

Policy regions

Value function with $\beta = 10$

General results: Defining "better" priors

- Definition: π_2 likelihood-ratio (LR) dominates π_1 ($\pi_2 \text{ LR } \pi_1$) if $\pi_2(\cdot) = \pi_1(\cdot)$ is increasing in \cdot .
- Examples of LR improvements:
 - Beta: Increasing β while holding the precision (τ) constant
 - Normal: Increasing the mean while holding the variance constant
- LR-dominance implies FOSD, but the reverse is not true.

Not a LR-improvement

A LR-improvement

- The LR-order survives Bayesian updating: given a signal

$$\pi_2 \text{ LR } \pi_1, \quad (\pi_2; X) \text{ LR } (\pi_1; X); \text{ for all } x \in X:$$

General results: Ordered signal processes

- Definition: The signal process $\{x_j\}$ satisfies the monotone-likelihood-ratio (MLR) property if the signal space \mathcal{X} is ordered and

$$L(x_j \geq 2) \text{ LR } L(x_j \geq 1) \text{ for all } 2 \geq 1 :$$

- Examples: Bernoulli signals; normal signals
- If the signal process satisfies the MLR property, then:

$$2 \text{ LR } 1) \quad f(x_2) \text{ LR } f(x_1)$$

For any prior π , $x_2 \geq x_1$, $(\pi; x_2) \text{ LR } (\pi; x_1)$:

General results: Increasing value functions

- Definition: $V(\cdot)$ is LR-increasing if $V(\omega_2) \geq V(\omega_1)$ whenever $\omega_2 \succ_{LR} \omega_1$.
- Proposition [Increasing]: Suppose the DM's utility function $u(w)$ is increasing in w and the signal process satisfies the MLR property. Then, for all k and w , the value function $U_k(w; \cdot)$ is LR-increasing in \cdot .

General results: Increasing value functions

- Value (or derived utility) function with k periods remaining:

$$\begin{aligned}
 U_0(w; \delta) &= u(w) && \text{(adopt)} \\
 &\approx E[u(w + \tilde{\tau})] \\
 U_k(w; \delta) &= \max_{\gamma} E[U_{k-1}(w - c; (\delta; \kappa))] && \text{(wait)} \\
 &u(w) && \text{(quit)}
 \end{aligned}$$

Value function with $\delta = 10$

General results: Increasing policies

- Proposition: Suppose the DM's utility function is increasing and the signal process satisfies the MLR property.

Rejection: If it is optimal to reject with prior π_2 , it is also optimal to reject with any prior π_1 such that $\pi_2 \succeq_{LR} \pi_1$.

Proof: Follows from LR-increasing value functions.

Adoption: Suppose the DM is risk neutral (or risk seeking). If it is optimal to adopt with prior π_1 , then it is also optimal to adopt with any prior π_2 such that $\pi_2 \succeq_{LR} \pi_1$.

Proof: Utility difference between adoption and waiting is LR-increasing.

- With risk neutrality, policies "increase" from quit to wait to adopt as LR-improves.

Illustrative example revisited

- Policies are LR-increasing.

Policy regions

Value function with $+ = 10$

Illustrative example revisited

- Policies are LR-increasing.

Policy regions

Value function with $+ = 10$

Illustrative example revisited


Risk neutral

Risk averse


- With risk aversion, can adopt and wait cross twice as we LR-improve ?

If the DM is risk averse, adoption policies may not be monotonic in .

- Example: log utility; three technology values; signals satisfy MLR property



Comparing waiting and adopting in this example:



- With log utility, a bad technology outcome + search costs can be catastrophic if the resulting wealth level is near zero.

Search costs push the DM “over the edge”

Can we ensure monotonicity by limiting the degree of risk aversion?

Increasing Adoption Policies

Proposition Suppose the DM is risk averse and her utility function u exhibits decreasing absolute risk aversion (DARA), i.e., her risk tolerance $u''(w)$ is increasing. Then, if

$$u''(w_0 + c) \leq u''(w_0)$$

where $w_0 = w - kc$ and $c = \min$ (“Not too risk averse”), then adoption policies are monotonic.

If u is CARA, no risk tolerance bound is required.

- We define a new property: “sLR-increasing”:

LR-increasing functions are sLR-increasing

sLR-increasing functions are single-crossing

Bayesian updating preserves sLR-increasing property

- We show utility difference between adoption and waiting is sLR-increasing

Then, utility difference between adoption and waiting is single crossing.

Generalizations:

- Model with discounting, $u(w + \text{NPV of costs/benefits})$

Delay is costly; also risk reducing

Results and proofs follow the same pattern

- Other applications of s-increasing: Monotonic policies in DPs

Submodularity (increasing differences) conditions (Topkis (1979), Lovejoy (1987a,b)) are sometimes hard to establish

Single-crossing conditions (e.g., Milgrom and Shannon (1994), Quah and Strulovici (2012), . . .) are hard to use in DPs

