

Behavioral Economics I

Thus far, we have assumed that an agent's utility is *quasilinear* and that an agent maximizes *expected utility*:

$$u_i(b_i, \mathbf{b}_{-i}) = x_i(b_i, \mathbf{b}_{-i})(v_i - p_i(b_i, \mathbf{b}_{-i})).$$

Note that this is equivalent to what we've been doing, where we've said that $u_i = x_i v_i - \hat{p}_i$, and here $p_i = \hat{p}_i / x_i$ when $x_i > 0$, and should otherwise be 0 anyway for IR.

However, humans do not always maximize this objective. Some environments and some individuals lead to different behavior. Consider the following.

Gamble A	\$10 w.p. 1
Gamble B	$\begin{cases} Y = \$20 & \text{w.p. } 0.5 \\ \$0 & \text{w.p. } 0.5 \end{cases}$

Which gamble do you prefer? For what value of Y would you prefer gamble B ?

What if the gambles instead look like the following:

Gamble A	\$1000 w.p. 1
Gamble B	$\begin{cases} Y + \$990 & \text{w.p. } 0.5 \\ \$990 & \text{w.p. } 0.5 \end{cases}$

How can we explain this behavior?

The Expected Utility Theory of von Neumann Morgenstern [15] says that this can be explained by agents having a non-linear value function for money.

$$u_i(b_i, \mathbf{b}_{-i}) = x_i(b_i, \mathbf{b}_{-i}) \cdot m(v_i - p_i(b_i, \mathbf{b}_{-i})).$$

One can arrive at the theory that agents are *expected utility* maximizers from a few small number of natural axioms:

- Completeness: For any two outcomes (including lotteries) A and B , an agent can rank these. (Either $A \succeq B$ or $B \succeq A$ or both.)
- Transitivity: If $A \succeq B$ and $B \succeq C$ then $A \succeq C$.

- Continuity: Let $A \succeq B \succeq C$. Then there exists some $p \in [0, 1]$ such that B is equally preferred to $pA + (1 - p)C$.
- Independence of Irrelevant Alternatives: If $A \succeq B$, then $tA + (1 - t)C \succeq tB + (1 - t)C$ for every lottery C and real $t \in [0, 1]$.

If these axioms are satisfied, then the agent is rational, and their preferences can be represented by maximizing expected utility w.r.t a utility function.

So what should $m(\cdot)$ look like in the above formulation of utility? EUT says that if $m(\cdot)$ is concave, that is, agents have *decreasing marginal returns* for money, then this explains risk-averse behavior. (And convex functions explain risk-seeking behavior.)

EUT Risk Aversion in Mechanism Design. There is a fair amount of work [8, 11, 12] in this area, as EUT is most commonly used model.

The Calibration Theorem (Rabin [16]):

Gamble A	$\begin{cases} \$11 & \text{w.p. } 0.5 \\ -\$10 & \text{w.p. } 0.5 \end{cases}$
Gamble B	$\begin{cases} \$Y & \text{w.p. } 0.5 \\ -\$100 & \text{w.p. } 0.5 \end{cases}$

Suppose an agent rejects gamble A, deciding it has negative utility for them. Under EUT, for what values of $\$Y$ will the agent also reject gamble B?

Answer: This rejection of gamble A bounds the curvature of a concave $m(\cdot)$ in a way such that the only possible answer is to reject gamble B for all values of $\$Y$, even if $\$Y = +\infty$.

Here's another experiment that was actually run, called the Allais Paradox [1]. Which gamble would you prefer?

Gamble A	$\$1\text{M w.p. } 1$
Gamble B	$\begin{cases} \$1\text{M} & \text{w.p. } 0.89 \\ \$5\text{M} & \text{w.p. } 0.10 \\ \$0 & \text{w.p. } 0.01 \end{cases}$

Most agents prefer gamble A, and EUT cannot explain this choice for any concave $m(\cdot)$. This motivated *cumulative prospect theory* (CPT) [10, 17]. CPT accounts for three things not seen in EUT:

- Probability weighting.
- Diminishing sensitivity.

- Loss aversion.

Probability weighting:

$$u_i(b_i, \mathbf{b}_{-i}) = y_i(x_i(b_i, \mathbf{b}_{-i})) \cdot m(v_i - p_i(b_i, \mathbf{b}_{-i})).$$

Suppose $y(\cdot)$ is convex, like x^2 . Then you value certain probabilities (0 and 1) the same, but you underweight uncertain probabilities.

Probability Weighting in Mechanism Design. Recent work [3] has examined revenue-optimal mechanism design with known probability weighting functions, unknown such risk attitudes, and in a dynamic setting as well.

Diminishing sensitivity is with respect to the reference point (see next). The closer to the reference point, the more sensitive.

Loss Aversion

Loss Aversion is the concept that losses weigh more heavily on an agent than gains. This requires evaluating gains and losses relative to some reference point.

An *alternative* (think: action) is a one-dimensional random variable c . Given some increasing utility function $m : \mathbb{R} \rightarrow \mathbb{R}$, faced with choice set C , an EU maximizer picks $c \in C$ to $\max \mathbb{E}[m(c)]$.

Incorporating loss aversion: Given a reference point $r \in \mathbb{R}$ and parameters $\eta \geq 0, \lambda \geq 1$:

$$n(c | r) = \begin{cases} \eta(m(c) - m(r)) & \text{if } m(c) - m(r) > 0 \\ \eta\lambda(m(c) - m(r)) & \text{otherwise} \end{cases}$$

Faced with choice set C , a loss-averse decision maker picks $c \in C$ to $\max \mathbb{E}[m(c) + n(c | r)]$. Then EUT is a special case. Think of $m(\cdot)$ as ‘material’ utility— m is the classical part, and n is the gain/loss part. This is a good compromise between accuracy and tractability.

Let’s use this to explain the difference in small/large stakes risk aversion. Suppose $m(c) = c$, $\eta = 1$, $\lambda = 1.5$. Let the reference point r refer to status quo wealth w . Then

$$m(w + x) + (w + x | r) = w + x + \begin{cases} x & \text{if } x > 0 \\ 1.5x & \text{otherwise} \end{cases}$$

The decision maker will accept the coin flip if and only if the ratio of gain to loss is at least 1.25.

The Endowment Effect (Kahneman Knetsch Thaler [9]): In a famous experiment, coffee mugs were randomly allocated to half of the class. The following (not truthful but can be made so) market-clearing mechanism was used. The students were asked for a list of prices: “At a price of $\$p$ will you sell/buy?” to which they answered yes or no. Then this was used to find a market-clearing price.

If we randomly allocate 22 mugs to 44 students, we should expect 11 trades. However, in practice, we only see about 1-5 trades.

What we observe is that there’s a gap between buyers’ stated willingness to pay and sellers’ willingness to sell.

One explanation given is that people have transaction costs: it’s a hassle to trade. However, if you use vouchers instead of mugs, you see exactly 11 trades.

It turns out that loss aversion with multiple dimensions can explain this. You have a reference point for the mug (whether you own it or not) and a reference point for money (how much you start with). You would model this as maximizing

$$\sum_i [m_k(c_k) + n_k(c_k | r_k)]$$

for a variety of reference points r_k .

The Endowment Effect in Mechanism Design. Recent work [2, 7] uses the endowment effect to show the existence of Walrasian Equilibria beyond Gross Substitutes valuations.

The Winner’s Curse and Cursed Equilibrium

Recall the *wallet game* with interdependent valuations. Each agent i knows the amount of money in their wallet, s_i , and their value is the amount of money in total in the room, $v = \sum_i s_i$. Whoever wins the auction (has the highest bid) for the sum of wallets gets all of the cash from the wallets. Let’s suppose we use an English (ascending) auction. For simplicity, suppose $s_i \sim U[0, 1]$. How should the players bid if they only know their own signal? Let’s suppose there are only two agents: you and your opponent.

What would you bid? What is $\mathbb{E}[s_i]$ for all i ?

Naive strategy: bid $s_i + \mathbb{E}[s_j]$, where $\mathbb{E}[s_j] = 0.5$.

BNE strategy: Quit at $2s_i$.

Crux: If you plan to quit at $2s_i$, then if you get to that bid, you know your opponent

has $s_j \geq s_i$.

Cursed equilibrium (Eyster and Rabin [6]) asserts: Players correctly understand the marginal distributions $\mathbf{b}_{-i} | s_i$ and $\mathbf{s}_{-i} | s_i$, but under appreciate how \mathbf{b}_{-i} and \mathbf{s}_{-i} are related.

It's as though with some probability χ , independent of \mathbf{s}_{-i} , all opponents play the unconditional distribution of \mathbf{b}_{-i} .

Let $\bar{\sigma}$ denote the strategy of your opponents averaged over their types. That is,

$$\bar{\sigma}_{-i}(\mathbf{b}_{-i} | s_i) = \sum_{\mathbf{s}_{-i}} \Pr[\mathbf{s}_{-i} | s_i] \cdot \sigma_{-i}(\mathbf{b}_{-i} | \mathbf{s}_{-i}).$$

Then σ is a χ -cursed equilibrium if for all i, θ_i : for all b_i such that $\sigma_i(b_i | s_i) > 0$:

$$b_i \in \arg \max_{b'_i} \sum_{\mathbf{s}_{-i}} \Pr[\mathbf{s}_{-i} | s_i] \sum_{\mathbf{b}_{-i}} u_i(b'_i, \mathbf{b}_{-i}, s_i, \mathbf{s}_{-i}) \cdot \left[\chi \bar{\sigma}_{-i}(\mathbf{b}_{-i} | s_i) + (1 - \chi) \sigma_{-i}(\mathbf{b}_{-i} | s_i) \right]$$

If $\chi = 0$, this is just a BNE. If $\chi = 1$, this is fully cursed—you operate as if everyone else is playing an average over all types and ignoring their true type.

Another example: Seller has value equal to their private signal s_S . Buyer has value $1.5s_S$. The seller will accept a price p if and only if $p \geq s_S$. What price p should the buyer offer? Say $s_S \sim U[0, 1]$.

The BNE has $p = 0$. This is not what we see in practice.

The Winner's Curse in Mechanism Design. Recent work [4] has examined designing individually rational but cursed EPIC welfare- and revenue-optimal mechanisms (second-best) and give a condition under which they approximate first-best welfare.

Other concepts: Quantal response [13, 14]: People choose actions randomly, but choose higher-utility outcomes with high probability.

Level- k Thinking [5]: We assume that each agent performs a finite number of steps of strategic reasoning.

- Level-0: Some default nonstrategic distribution of play (perhaps uniform).
- Level-1: Best respond to level-0 players.
- ...
- Level- k : Best respond to level $k - 1$, or possibly levels $\{0, 1, \dots, k - 1\}$.

Acknowledgements

This lecture was developed by using materials from Shengwu Li, and in particular, his 2019 tutorial on Behavioral Economics and Imperfect Rationality at ACM EC.

References

- [1] Maurice Allais. Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica: journal of the Econometric Society*, pages 503–546, 1953.
- [2] Moshe Babaioff, Shahar Dobzinski, and Sigal Oren. Combinatorial auctions with endowment effect. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 73–90, 2018.
- [3] Shuchi Chawla, Kira Goldner, J. Benjamin Miller, and Emmanouil Pountourakis. Revenue maximization with an uncertainty-averse buyer. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '18*, pages 2050–2068, USA, 2018. Society for Industrial and Applied Mathematics.
- [4] Yiling Chen, Alon Eden, and Juntao Wang. Cursed yet satisfied agents. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.
- [5] Vincent P Crawford and SESS Distinguished Lecture. Outguessing and deception in novel strategic situations. *School of Economics and Social Sciences Distinguished Lecture, Singapore Management University*, 2004.
- [6] Erik Eyster and Matthew Rabin. Cursed equilibrium. *Econometrica*, 73(5):1623–1672, 2005.
- [7] Tomer Ezra, Michal Feldman, and Ophir Friedler. A general framework for endowment effects in combinatorial markets. *ACM SIGecom Exchanges*, 18(2):38–44, 2020.
- [8] Audrey Hu, Steven A Matthews, and Liang Zou. Risk aversion and optimal reserve prices in first-and second-price auctions. *Journal of Economic Theory*, 145(3):1188–1202, 2010.
- [9] Daniel Kahneman, Jack L Knetsch, and Richard H Thaler. Experimental tests of the endowment effect and the coase theorem. *Journal of political Economy*, 98(6):1325–1348, 1990.
- [10] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.

- [11] Eric Maskin and John Riley. Optimal auctions with risk averse buyers. *Econometrica: Journal of the Econometric Society*, pages 1473–1518, 1984.
- [12] Steven A Matthews. Selling to risk averse buyers with unobservable tastes. Technical report, Discussion Paper, 1981.
- [13] Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995.
- [14] Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for extensive form games. *Experimental economics*, 1:9–41, 1998.
- [15] J von Neumann, Oskar Morgenstern, et al. Theory of games and economic behavior, 1947.
- [16] Matthew Rabin. Diminishing marginal utility of wealth cannot explain risk aversion. 2000.
- [17] Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5:297–323, 1992.