Entropy-constrained behavior and the endowment effect

by Duncan K. Foley

Department of Economics, New School for Social Research, 79 Fifth Avenue, New York, NY 10003, foleyd@newschool.edu.

Acknowledgments

I'd like to thank Ali Khan for stimulating conversations that motivated this work, and Nicholas Foley, Jang-ho Yang, Ellis Scharfenaker, and Gregor Semieniuk for their useful comments. These slides were prepared for the Columbia Conference Capitalism, Conflict, and Cooperation: A Celebration of the Work of Samuel Bowles and Herbert Gintis on October 5, 2019. The interested reader can find a more detailed discussion of the topics in the working paper with the same title available at https://EconPapers.repec.org/RePEc:new:wpaper:1731. My interest in the endowment effect traces to many conversations with Sam Bowles.

Choice behavior

- The conventional economic theory derived from marginalist thinking universally assumes that individual economic agents have complete and consistent preferences over outcomes and behave to maximize preferences.
- This axiom, however, has implications that are inconsistent with widely-observed human (and, in fact, animal) behavior, and leads to a series of paradoxes and mathematical complications in explaining economic phenomena.

Basic framework

- Consider a subject who faces an offer to sell an asset the subject values at μ in terms of money for a price p.
- The decision-maker, according to the assumption of complete and consistent preferences, either prefers the money price to the asset, in which case she accepts the offer, or prefers the asset to the offered money price, in which case she refuses the offer.
- As illustrated in Figure 1 this implies a sharp step function in the frequency with which the decision-maker accepts the offer as the price changes.
- The price at which the decision-maker shifts her behavior can in principle be determined to any desired degree of precision.

Visualizing choice behavior

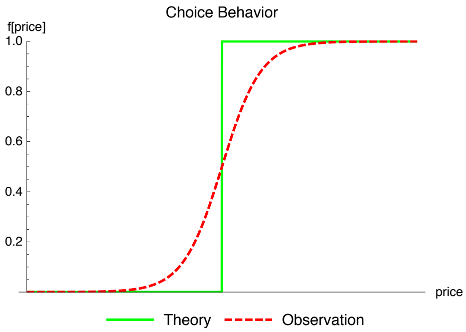

Figure 1: Conventional choice theory predicts a sharp step-function response in the frequency with which subjects accept an offer to exchange some good for money as the price changes (green curve). Observations invariably show a logistic quantal response to changes in price (red dashed curve), governed by a parameter that determines how close the response is to the step function.

Quantal response

- One of the best-confirmed results of quantitative psychology and experimental social science, however, is *quantal response*, partial randomization of the responses of human beings (and other organisms) to environmental stimuli as the stimuli move through regions of transition from one stimulus to another.
- In the context of the choice described here, this implies behavior according to the red dashed curve in Figure 1.
- The slope of the quantal response is determined by a parameter which I will call the *behavior temperature*, T, and varies from subject to subject and context to context.

Natura non facit saltum

- The sharp step-function behavior predicted by expected utility theory (and other economic theories of choice that assume behavior reflects preference maximization over complete and consistent preferences) is at odds with the presumption of natural scientists that, except in extreme circumstances of only theoretical significance, like temperatures of absolute zero, systems tend to exhibit the smoothest behavior compatible with the constraints imposed by natural laws.
- In mathematical terms this principle translates into the rule that informational entropy cannot be arbitrarily low. The informational Shannon entropy of a frequency distribution $\{f_1, \ldots, f_K\}$ with $f_k \geq 0, \forall k, \sum_k f_k = 1$ is $H = -\sum_k f_k \mathrm{Log}\left[f_k\right]$ with the convention that $0\mathrm{Log}[0] = 0$.

Mixed strategy choice behavior

- In the context of choice theory, consider the problem of a decision-maker who can choose between actions, say sell or not sell, $a_1, \ldots, a_k$, knowing the payoff $u[a_k]$ for each.[1]
- In general we can represent the behavior of the decision-maker as a mixed strategy assigning some non-negative frequency $f_k \geq 0$ to each of the actions, and resulting in an expected payoff $\sum_k f_k u[a_k]$.
- Maximizing expected utility subject only to the normalization constraint $\sum_k f_k = 1$ will lead the decision-maker to choose the highest-payoff action with frequency 1 and all the others with frequency 0. The entropy of this distribution is zero.

[1] The term "utility" carries with it both the sense of "payoff" and the sense of "welfare". Since welfare is not the issue in this context, I will use the term "payoff" to describe the variable that influences choice behavior. Conversations with Sam Bowles alerted me to this point.

Entropy-constrained mixed-strategy choice behavior

- If, however, the decision-maker faces a lower bound on the entropy of her mixed strategy, $\bar{H}$, she solves the mathematical programming problem:

$$\underset{\{f_1,\ldots,f_K\}\geq 0}{\text{Max}} \sum f_k u[a_k] \text{ subject to} \sum f_k = 1, -\sum f_k \text{Log}[f_k] \geq \bar{H} \quad (2)$$

- Because the objective function of this program is linear in the frequencies, and the constraint defines a convex set, the first-order conditions of the Lagrangian

$$\mathcal{L}[f, \mu, T] \quad (3)$$

$$= \sum f_k u[a_k] - \mu\left(\sum_k f_k - 1\right) + T\left(-\sum f_k \text{Log}[f_k] - \bar{H}\right)$$

are necessary and sufficient to characterize the solution.

Gibbs distribution

- The first-order conditions can be solved to yield:

$$f\left[a_k\right] = \frac{e^{\frac{u[a_k]}{T}}}{\sum_k e^{\frac{u[a_k]}{T}}} \tag{4}$$

- This is the *Gibbs distribution*, which leads the decision-maker to choose each available action with a positive frequency, with the logarithm of frequency proportional to the ratio of the payoff to the Lagrange multiplier T.
- In physical systems T, the Lagrange multiplier corresponding to entropy is referred to as a *temperature*, and in the behavioral context can be regarded as a *behavioral temperature*.
- The lower the behavioral temperature, the more concentrated the decision-maker's behavior is on the payoff-maximizing action.

Quantal response is a Gibbs distribution

- In the case where there are only two actions, the Gibbs distribution is

$$f[a_1] = \frac{e^{\frac{u[a_1]}{T}}}{e^{\frac{u[a_1]}{T}} + e^{\frac{u[a_2]}{T}}} = \frac{1}{1 + e^{\frac{u[a_2]-u[a_1]}{T}}} \quad (5)$$

- This is the logistic function experiments reveal
- One way (though certainly not the only way) to understand logistic behavior is reflecting informational constraints on human responses to stimuli such as a choice situation, that is, as *entropy-constrained behavior*.[2]
- The same distribution results from maximizing entropy subject to a constraint on expected payoff as in models of *satisficing*.

[2] This derivation of logistic quantal response behavior is essentially equivalent to Christopher Sims' theory of "rational inattention" (Sims, 2012).

Quantal response and choice theory

- Logistic quantal response behavior is a generalization of rational choice theory, in so far as the decision-maker, as in rational choice theory, has well-defined payoffs over actions, and maximizes expected utility in choosing a mixed strategy, which results in more frequent choices of higher-payoff actions.
- The new element in entropy-constrained behavior is the behavior temperature, which limits the degree to which the decision-maker can concentrate frequency on the highest-payoff action.
- The implications of entropy-constrained logistic behavior are far-reaching, but much research in psychology and economics is conducted and interpreted without taking these implications fully into account.

Entropy-constraints and preferences

- Logistic quantal response behavior violates the assumptions of consistency and completeness of preferences
- This violation arises because subjects behaving according to a logistic quantal response sometimes choose one option and sometimes another when presented with exactly the same choice.
- Conventional economic theory has responded to this anomaly by seeking one or another way to defend the assumptions of consistency and completeness of preferences and the principle that observed choice reflects (unconstrained) payoff maximization.

Qualitative effects of entropy constraints

- It is tempting to think that because entropy-constrained logistic quantal response behavior approximates full payoff-maximizing behavior the converse is true, so that models assuming full payoff-maximizing behavior are reliable guides to entropy-constrained behavior in the real world.
- But this logic does not hold, because the limiting case where behavior temperature goes to zero has important qualitative differences from entropy-constrained behavior at any positive behavior temperature.
- For example, at any positive behavior temperature the entropy-constrained model predicts that we will observe every available action with some positive (though possible very low) frequency. But unconstrained payoff-maximization predicts that we will observe only payoff-maximizing actions.

Entropy constraints and conventional choice theory

- This point touches the most fundamental aspects of conventional economic theory, including the analysis of market equilibrium through supply and demand curves, endowment and other loss-aversion effects, the interpretation of Cournot-Nash equilibria in games, the distributional effects of market exchange, and such key concepts as Bertrand's "cut-throat competition".
- The problem is that many conclusions that hold for the knife-edge case $T = 0$ do not hold qualitatively in the case $T > 0$, no matter how low the behavior temperature is.

The endowment effect

- One important implication of entropy-constrained choice behavior involves observations on a population of subjects even when they are not actively engaged in economic interactions, and they all have identical payoffs and behavior temperatures.
- Amos Tversky and Daniel Kahneman Tversky and Kahneman (1991); Kahneman et al. (1991) demonstrated the experimental replicability of a wide range of subject behaviors that are anomalous from the point of view of rational choice theory under the general rubric of "loss-aversion".
- One widely-noticed instance is the "endowment effect", which purports to demonstrate that ownership of some good changes the payoff associated with it to is owner.
- It is instructive to consider the endowment effect from the point of view of entropy-constrained decision theory.

The endowment effect scenario

- In one highly influential and often-replicated experiment, a group of subjects (typically a college class) are randomly divided into two sub-groups, one of which receives an "endowment" of an object (such as a coffee mug) of moderate value.
- The subjects then report the prices at which they would sell or buy the object.
- The mean selling prices reported by the sub-group that received the object are replicably and statistically significantly higher than the mean buying prices reported by the sub-group who did not receive the object.
- Kahneman and Tversky interpret these reported subjective prices as point estimates of the subjects' payoffs, and regard the difference as an *endowment effect*, reflecting a context-dependent attachment of the subjects to the things they own.

Endowment effect might be an illusion

- The observed difference in reported buying and selling prices, however, is also a direct implication of entropy-constrained behavior, without assuming changed payoffs as a result of ownership
- Entropy-constrained behavior theory predicts subjects offered the opportunity to buy or sell at various prices will respond along a logistic curve at some non-zero behavior temperature
- At very low prices the subject would (almost) always buy, and at very high prices (almost) always sell, but for some intermediate range of prices the subject will buy with a frequency that falls with the offered price.
- When asked at what price they would buy or sell, subjects may interpret the question in terms of a frequency threshold and report the price at which they would buy or sell the object with that frequency.

Visualizing the endowment effect

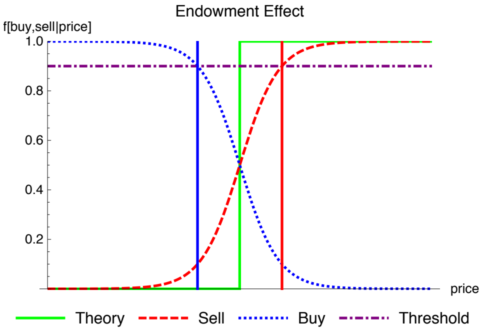

Figure 2: An entropy-constrained economic agent will buy or sell a good with a frequency depending on the price offered (dashed red and dotted blue curves) given her value (the vertical section of the green curve). The agent reports behavior at some frequency threshold (the dashed-dotted purple curve), implying a lower buying than selling price, despite having an unchanged valuation of the good.

Replicability of the endowment effect

- The issue here is not the empirical replicability of experimental observations.
- Logistic quantal response behavior occurs reliably in a wide range of choice situations, and there is no reason not to believe that the experimental data reported as supporting loss aversion is highly replicable.
- The issue is the interpretation of these experiments as indicating a shift in underlying agent payoffs due to context.

Interpretation of the endowment effect

- Conventional choice theory, interpreted as the existence of consistent and complete preferences and the principle that actual behavior represents unconstrained maximization of expected payoff, is surely inconsistent with the evidence from the endowment effect experiments
- The question is what modifications in theory are called for by these anomalies
- Entropy-constrained quantal response behavior generalizes conventional choice theory through the introduction of a single new parameter, the behavior temperature
- The theory of loss-aversion, on the other hand, requires the introduction of new parameters for every situation