

Information theory and behavior

Duncan K. Foley^a

Department of Economics, New School for Social Research, 79 Fifth Avenue, New York, NY 10003, USA

Received 8 July 2019

Published online 7 July 2020

Abstract. The quantal response behavior widely observed in experiments and observations of human and animal behavior can be derived as expected payoff maximization subject to a constraint on the entropy of the subject's behavior mixed strategy. The Lagrange multiplier corresponding to the entropy constraint is an agent's "behavior temperature". Entropy-constrained behavior approximates payoff-maximizing behavior, but in many contexts exhibits qualitatively different outcomes. The "endowment effect" and other instances of "loss-aversion", for example, can be seen as a consequence of entropy-constrained behavior. Identical entropy-constrained agents with the same value for a good or asset will exhibit spontaneous "noise trading". An entropy-constrained agent with a lower behavior temperature will systematically take economic surplus away from an agent with the same valuation of a good but a higher behavior temperature in bilateral transactions. The equilibrium of a standard supply-demand models with entropy-constrained agents is a non-degenerate frequency distribution of transaction prices rather than a single equilibrium price. Changes in behavior temperature can transform social interaction games from prisoners' dilemmas to assurance games. Entropy-constrained quantal responses allow quantitative inferences about payoff changes and distribution stronger than qualitative Pareto comparisons.

1 Choice behavior

The conventional¹ 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. It is my purpose in this paper to explore these points in detail.

Consider a decision-maker who holds a bundle of goods (x_1, \dots, x_K) confronted with an offer to receive a quantity of good k , Δx_k , in exchange for a given quantity of the first good, which we treat as the numéraire, $-\Delta x_1$ (the *price*). The decision-maker, according to the assumption of complete and consistent preferences, either

^a e-mail: foleyd@newschool.edu

¹ The results in Section 4 come from conversations with Nicholas C. Foley. I'd like to thank Ali Khan for stimulating conversations that motivated this work, and Jang-ho Yang, Ellis Scharfenaker, and Gregor Semieniuk for their useful comments.




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

prefers the proposed bundle $(x_1 - \Delta x_1, \dots, x_k + \Delta x_k)$ to her original bundle, in which case she accepts the offer, or prefers her original bundle to the proposed bundle, 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. This implies that the price at which the decision-maker shifts her behavior can in principle be determined to any desired degree of precision.

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, and varies from subject to subject and context to context.

A leading and paradigmatic economic example is the finding of Duncan Luce in experiments on human subjects' risk aversion. In these experiments subjects were offered the choice between a prize B with certainty, and a lottery offering with frequency $\frac{1}{2}$ a prize A that the subject consistently preferred to B when offered each with certainty, and with frequency $\frac{1}{2}$ a prize C that the subject consistently preferred B to when offered each with certainty. The *expected utility* theorem of John von Neumann and Oskar Morgenstern predicts that an expected utility maximizer has payoffs for the prizes $u[A], u[B], u[C]$, measurable up to an affine transformation $u' = a + bu$ and values lotteries according to the mathematical expectation of their payoffs. This theory implies there is a unique outcome payoff $u^* = \frac{1}{2}u[A] + \frac{1}{2}u[C]$, such that the subject will always choose the certain prize when $u[B] > u^*$, and always choose the lottery when $u[B] < u^*$. (Because payoffs are determined only up to an affine transformation, without loss of generality we can take $u[C] = 0, u[A] = 1$.) With this payoff scale the subject's certainty equivalent payoff $u^* = \frac{1}{2}u[A]$. Luce discovered, in line with earlier and independent contemporary psychological researchers, that it is impossible to determine the payoffs with arbitrary precision because in any experimental setting subjects randomize their choices when the certainty equivalent payoff

is close to u^* . When $u[B]$ is sufficiently high subjects will (almost) always choose the certain prize, and when $u[B]$ is sufficiently low subjects will (almost) always choose the lottery, but in an intermediate interval that cannot be completely eliminated, subjects sometimes choose the lottery and sometimes choose the certain prize.

Luce found, moreover, that the frequency with which subjects chose the lottery, f , in this type of experimental setting consistently followed the logistic law

$$f = \frac{e^{\frac{\frac{1}{2}u[A]-u[B]}{T}}}{1 + e^{\frac{\frac{1}{2}u[A]-u[B]}{T}}}$$

for a parameter $T > 0$ that varied with the subject and the experimental protocol.

Although Luce’s experiments involved choices over lotteries, the quantal response also describes choice behavior between certain outcomes.

Figure 1 summarizes the situation. The horizontal axis represents the price for some action presented to the subject, and the vertical axis the frequency with which the subject takes the action given that price (the probability p of success in a series of Bernoulli trials) While expected utility theory predicts a sharp step-function response of an expected utility-maximizing subject to changes in the certain payoff, experimental evidence invariably finds an S-shaped logistic response. The closeness of the logistic response to the step function is expressed by the parameter T . In the (empirically unreachable) limit $T \rightarrow 0$, the quantal response converges to the (Heaviside) step function.

1.1 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 $\{f_1, \dots, f_K\}$ with $f_k \geq 0, \forall k, \sum_k f_k = 1$ is $H = -\sum_k f_k \text{Log}[f_k]$ with the convention that $0\text{Log}[0] = 0$.

In the context of choice theory, consider the problem of a decision-maker who can choose among a finite number of actions a_1, \dots, a_K , knowing the payoff $u[a_k]$ for each². 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 (assuming we perturb payoffs so as to avoid ties). The entropy of this distribution is zero. If, however, the decision-maker faces a lower bound on the entropy of her mixed strategy, \bar{H} , she solves the mathematical programming problem:

$$\text{Max}_{\{f_1, \dots, f_K\} \geq 0} \sum f_k u[a_k] \text{ subject to } \sum f_k = 1, -\sum f_k \text{Log}[f_k] \geq \bar{H}.$$

Because the objective function of this program is linear in the frequencies, and the entropy is a strictly concave function of the frequencies that defines a convex set for

² 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.

the constraints, the first-order conditions of the Lagrangian $\mathcal{L}[f, \mu, T] = \sum f_k u[a_k] - \mu (\sum_k f_k - 1) + T (-\sum f_k \text{Log}[f_k] - \bar{H})$ are necessary and sufficient to characterize the solution. The first-order conditions can be solved to yield:

$$f[a_k] = \frac{e^{\frac{u[a_k]}{T}}}{\sum_k e^{\frac{u[a_k]}{T}}}. \quad (1)$$

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.

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

$$f[a_1] = \frac{e^{\frac{u[a_1]}{T}}}{e^{\frac{u[a_1]}{T}} + e^{\frac{u[a_2]}{T}}} = \frac{1}{1 + e^{\frac{u[a_2] - u[a_1]}{T}}}.$$

This is the logistic function psychological experimentation reveals. One way (though certainly not the only way) to understand logistic behavior is to regard it as reflecting informational constraints on human (or more generally organismic) responses to stimuli such as a choice situation, that is, as *entropy-constrained behavior*³.

Logistic quantal response behavior can be regarded as 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.

While logistic quantal response behavior can be regarded as a kind of generalization of conventional choice theory, it 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. For example, Luce himself derived the logistic quantal response function he observed empirically by assuming that the subject is uncertain about the payoff, with the uncertainty represented by a particular frequency distribution that implies logistic quantal response behavior. Later work in this field by economists and econometricians such as Charles Manski and Daniel McFadden [2] follows this line of interpretation. Unfortunately, as I will argue here, this way of rationalizing logistic quantal response behavior obscures important economic implications and can lead to misleading conclusions. One example is discussed in Section 2.

³ This derivation of logistic quantal response behavior is essentially equivalent to Christopher Sims' theory of "rational inattention" [5].

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.

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 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.

Entropy-constrained behavior leads to smooth average demand functions, even when the choice set, preferences, and constraints do not meet the convexity requirements for continuous demand functions with unconstrained maximization. This feature of entropy-constraints may help address some of the perplexities produced in general equilibrium theory because of its commitment to payoff maximization without entropy constraints⁴.

2 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 [1,6] 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 its owner. It is instructive to consider the endowment effect from the point of view of entropy-constrained decision theory.

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.

The difference in buying and selling prices, however, is also a direct implication of entropy-constrained behavior. If these subjects were offered the opportunity to buy or sell the object at various prices, entropy-constrained behavior theory predicts they would respond along a logistic curve at some non-zero temperature. At very low prices the subject would (almost) always buy, and at very high prices (almost) always

⁴ For example [3].




Fig. 2. An entropy-constrained economic agent whose valuation of some good is represented by the vertical section of the green curve will buy or sell the good with a frequency depending on the price offered, as represented by the dashed red “Sell” and dotted purple “Buy” logistic quantal response curves. If asked at what price she would buy or sell the good, the agent chooses some frequency threshold (the dashed-dotted purple curve) to represent a summary of her behavior, and will always report a lower buying price than selling price, despite having a well-defined and stable valuation of the good.

sell, but for some intermediate range of prices the subject will buy with a frequency that falls with the offered price. If the subjects in the endowment experiment behave in this way, it is reasonable to suppose that when asked at what price they would buy or sell the object they interpret the question as referring to some frequency threshold, such as 90%, and report the price at which they would buy or sell the object with 90% frequency.

As Figure 2 illustrates, if behavior temperature is non-zero, a typical subject will report a gap between buying and selling prices in this situation, despite having only a single unchanged payoff for acquiring or keeping the object. The price at which the object is offered or which is offered for the object is plotted on the horizontal axis, and the frequency with which a typical subject will buy (the blue dotted curve) or sell (the red dashed curve) is plotted on the vertical axis. The threshold frequency the subjects interpret as “willingness to buy or sell” is the purple dashed-dotted line. All the subjects have the same payoff for owning the object, but the theory of the endowment effect interprets the differences in reported prices for buying and selling as representing a shift in a single (zero behavior temperature) valuation of the object as a result of one sub-group owning it. In the constrained-entropy interpretation of the experiment, the difference is due to the fact that the sub-group who own the object are asked about their selling price and the sub-group who do not own the object are asked about their buying price. But researchers who firmly adhere to the principle that buying and selling behavior reflect exact maximization of an underlying payoff will interpret this difference as a change in the agent’s underlying payoff for the good depending on context. The reader can work out how other instances of apparent loss-aversion, possibly more disguised, will arise from entropy-constrained behavior.

The issue here is certainly 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 are highly replicable. The issue is the interpretation of these experiments as indicating a shift in underlying agent payoffs due to context.

Conventional choice theory, interpreted as the existence of consistent and complete preferences that can be represented by payoffs, and the principle that actual behavior represents unconstrained maximization of expected payoff, is surely inconsistent with the evidence from the endowment effect experiences. The question researchers face is what modifications in the theory are called for by these anomalies. Recognizing that entropy-constrained quantal response behavior is ubiquitous generalizes conventional choice theory through the introduction of a single new parameter, the behavior temperature. The theory of loss-aversion, on the other hand, practically requires the introduction of new parameters for every experimental and observational situation.

This is not to deny that even the generalized entropy-constrained version of choice theory may be flawed as a tool for the analysis of particular choice situations, and therefore in need of modifications beyond the introduction of entropy constraints.

3 Spontaneous transactions in markets

A striking example of the failure of conclusions that hold in the limit $T \rightarrow 0$ to be robust qualitatively when $T > 0$ is the question of whether identical agents will actually transact. In the case where agents are payoff-maximizers without an entropy constraint, identical agents have no incentive to transact because there are no potential gains from trade between them.

The case of entropy-constrained agents is, however, qualitatively different. Suppose the typical agent values the good at μ , has a payoff of buying (selling) a unit of the good at price p equal to $u[p] = \mu - p$, which is her consumer surplus when her valuation of the good is μ and operates at a behavior temperature $T > 0$. Then the frequencies of buying, $f[p]$, and selling, $1 - f[p]$, at price p are:

$$f[p] = \frac{1}{1 + \exp[-\frac{\mu-p}{T}]} \quad (2)$$

$$1 - f[p] = \frac{1}{1 + \exp[\frac{\mu-p}{T}]} \quad (3)$$

The frequency of a transaction at price p is $f[p](1 - f[p])$, the frequency with which one agent buys at that price and a counterpart agent sells. The frequency of transactions at price p , $\tau[p]$ is, assuming $\mu = 0$, so that p represents the difference between the transaction price and the typical agent's value of the good:

$$\tau[p] = \frac{1}{1 + \exp[\frac{p}{T}]} \frac{1}{1 + \exp[-\frac{p}{T}]} = \frac{1}{2 + 2 \cosh[\frac{p}{T}]}.$$

Integrating the frequency over all prices shows that the frequency of spontaneous transactions among entropy-constrained transactors who have the same valuation of a good or asset and operate is equal to their behavior temperature, T .

Entropy-constrained behavior provides an explanation for the widely-observed phenomenon of "noise-trading" in asset and other markets. The derivation shows that this phenomenon is closely connected to the endowment effect.

4 Distribution with different behavior temperatures

Even if agents' valuations of a good or asset are the same, differences in behavior temperatures will lead to unequal gains from transactions, as Figure 4 shows.

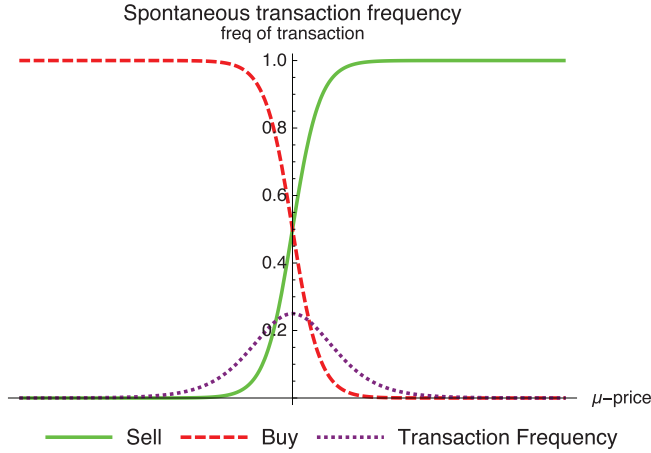


Fig. 3. Identical transactors with the same valuation of a good or asset and the same positive behavior temperature will generate non-zero spontaneous transactions due to quantal response effects. The frequency of transactions is equal to T , the area under the frequency curve.




Fig. 4. When agents with the same valuations but different behavior temperatures transact, the agent with the lower behavior temperature (the red quantal response) will take money away from the agent with the higher behavior temperature (the green quantal response). The high temperature agent sells more often at a low price, and buys more often at a high price. The net gain is represented by the area between the curves (remembering that the buyer gains at low prices and the seller at high prices), equal to $(T_{hi} - T_{lo}) \ln 2$.

With entropy-constrained buying and selling behavior, an agent with a higher behavior temperature sells more often at lower prices, and buys more often at higher prices. An example is the ability of experienced bond traders (who can value bonds relatively precisely) to profit from trades with less experienced customers (who can value the bonds relatively less precisely). The net transfer of wealth when the payoff is the difference between transaction price and the transactor’s valuation of the good is the area between the two quantal responses, $(T_{hi} - T_{lo}) \ln 2$, and is thus directly proportional to differences in behavior temperature.

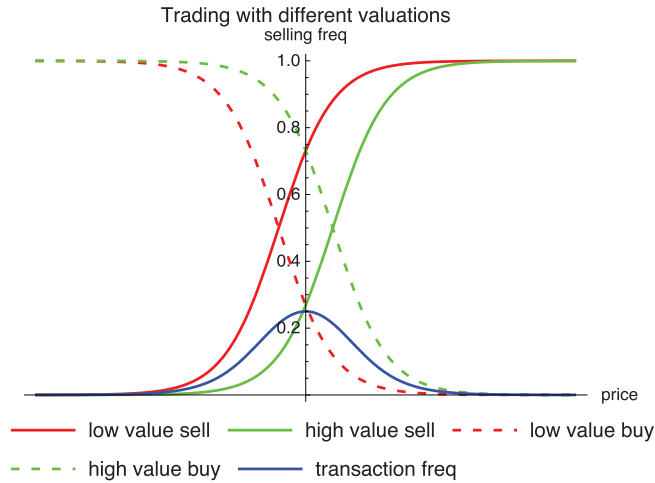


Fig. 5. When entropy-constrained agents have different valuations of the good or asset, transactions arise both spontaneously and as a means of realizing potential economic surplus. The figure plots selling and buying quantal responses for agents with a low valuation of the good in red and for agents with a high valuation of the good in green. The resulting frequency of transactions at different prices is plotted in blue. This interaction approximates supply-demand equilibrium (given the assumption of equal numbers of high- and low-valuation agents) in that the modal and mean transaction price is the average of the agents' valuations of the good.

This is another example of a significant qualitative difference between assuming that agents maximize payoffs without an entropy constraint, and assuming that there is an entropy constraint leading to a non-zero behavior temperature. Since the traders in this example have the same valuation of the good, if they have zero behavior temperatures, there will be no transactions at all, similarly to Figure 3. The transactions in this scenario neither realize nor destroy economic surplus, since the agents have the same valuation of the good, but serve only to transfer economic surplus from high to low behavior temperature transactors.

5 Transactions with different valuations

Conventional economic theory, based on the assumptions that behavior represents maximization of consistent and complete preferences without an entropy constraint, generally focuses on cases where there are differences between agents' valuations of goods (their marginal rates of substitution or marginal costs), so that there is an opportunity for mutually advantageous transactions. Figure 5 shows the frequency of transactions at various prices in this scenario.

In this case there are four possible combinations of agents that can give rise to transactions: a high valuation agent selling to another high valuation agent; a low valuation agent selling to another low valuation agent; a high valuation agent selling to a low valuation agent; and a low valuation agent selling to a high valuation agent. In the general case these possible configurations have to be weighted by the relative number of agents of each type.

In this scenario some transactions actually reduce the realization of economic surplus. This is an unavoidable consequence of the fact that agents maximize payoff subject to an entropy constraint, and therefore sometimes make what conventional economic theory would regard as erroneous transactions.

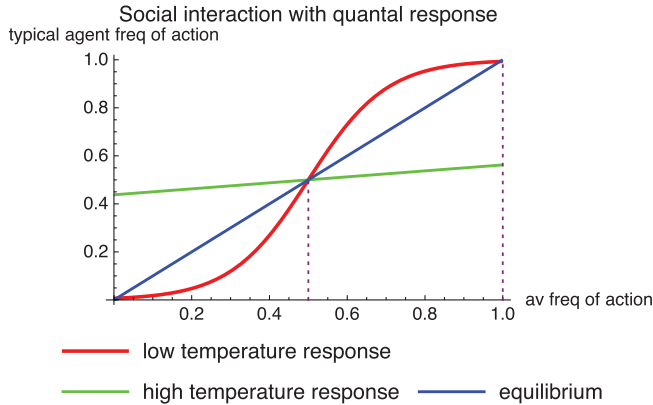


Fig. 6. When the frequency with which other identical agents choose some action influences the frequency with which a typical agent chooses that action, quantal response behavior temperature determines the number and stability of equilibria (which occur on the 45° line where all agents act identically). The horizontal axis measures the frequency with which other agents choose the action and vertical axis the frequency with which the typical agent’s best response is to choose the action, represented by two quantal response curves. The green quantal response represents a high behavior temperature: as a result the typical agent does not respond very much to the behavior of other agents, and there is a single stable interior equilibrium. The red quantal response represents a low behavior temperature: as a result the typical agent responds sensitively to the behavior of other agents and the interior equilibrium becomes unstable and bifurcates into two stable extreme equilibria. The behavior temperature can transform an interaction from a Prisoners’ Dilemma-like single equilibrium interaction into an Assurance Game-like multiple equilibrium interaction with path dependency.

6 Social interaction

Another important consequence of behavior temperature is the multiplicity and stability of equilibria in a social interaction scenario.

In this scenario the agent responds not to an offered price, but to the average behavior of other identical agents, measured by the frequency of their choosing some action. Because the agents are identical, equilibria occur where the quantal response frequency of the typical agent is equal to the frequency of the other agents, represented by the 45° line in Figure 6. When the behavior temperature of the typical agent is high, she does not respond much to other agents’ actions, and there is a single stable interior equilibrium. Given the strategic complementarity implied by the upward-sloping quantal best response function, this equilibrium will in general be a Prisoners’ Dilemma-like outcome with agents choosing the action too rarely. When the behavior temperature of the typical agent is low enough, her quantal best response cuts the equilibrium locus from below at the interior equilibrium, which becomes unstable, bifurcating into two stable extreme equilibria. In typical situations, one or the other of the stable equilibria is preferred by the typical agent, but which one prevails depends on the initial starting point so that the system is a path-dependent Assurance Game-like interaction. A change in behavior temperature transforms the Prisoners’ Dilemma-like interaction into an Assurance Game-like interaction⁵.

The quantal response of the typical agent induces a Markov chain on the state space of profiles of agent behavior in the social interaction model. Given any starting

⁵ For the stability of social interactions and see, for example [4].

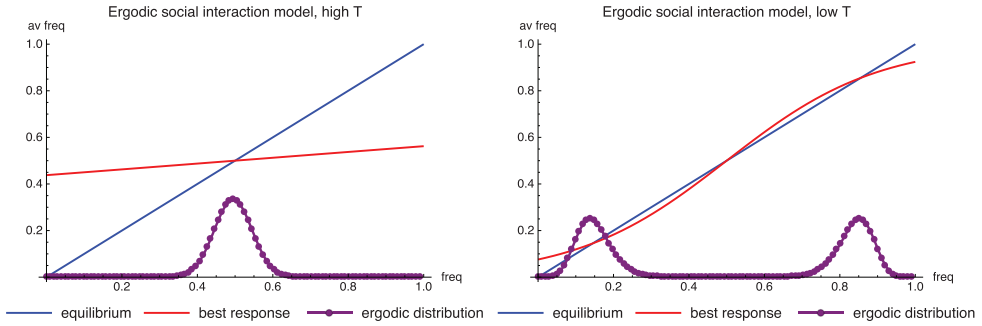


Fig. 7. Ergodic distributions of the social interaction model at high and low behavior temperatures. The red line is the best response, which is flat for high behavior temperature, and steep for low behavior temperature. The blue line is the equilibrium locus. The purple distribution dots represent the ergodic distribution induced by the social interaction in a population of $n = 100$. For high behavior temperatures the ergodic distribution is centered on the interior stable equilibrium. For low behavior temperatures the ergodic distribution is bi-modal concentrated on extreme outcomes in which most of the agents choose either to take or not to take the action. In the low behavior temperature case there is a positive frequency of transition between the extreme quasi-equilibrium configurations.

profile of agent actions determining the average frequency of the action, the quantal best response of the typical agent determines the frequency with which any agent chooses the action in the next period, and, as a consequence, the frequency of any particular profile of agent actions in the next period. (In the binary case this is just a binomial distribution.) If there are n agents each with behavior temperature T and a payoff $\mu - P$, where P is the average frequency of agents choosing the action, the state of the system in any round of interaction can be described as the number of agents choosing the action, $k = 0, \dots, n$, the average frequency of taking the action will be $P = \frac{k}{n}$, the frequency with which each agent will choose the action is $f[P] = \frac{1}{1 + \exp[\frac{\mu - P}{T}]}$ and the transition probabilities from state k to state k' , $t_{k,k'}$ are:

$$t_{k,k'} = \binom{n}{k'} f\left[\frac{k}{n}\right]^{k'} (1 - f\left[\frac{k}{n}\right])^{n-k'}. \tag{4}$$

The resulting Markov chain has an ergodic distribution, which describes the long run evolution of the system, as Figure 7 illustrates.

The ergodic distribution for high temperature social interactions is centered on the interior stable equilibrium, but there is a finite frequency for any outcome due to the fact that the quantal response has a non-zero entropy. The ergodic distribution for low temperature social interactions is bi-modal, with concentrations at quasi-equilibria corresponding to the extreme stable equilibria. Again, due to the fact that the quantal response has non-zero entropy, there is a finite frequency of transition between the extreme quasi-equilibria⁶.

7 Quantitative welfare economics

As is the case for other applications of expected utility theory, entropy-constrained quantal response implies that payoffs are cardinal up to an affine transformation of

⁶ For the dynamics of transitions between multiple quasi-equilibria in Markov chains see, for example [7].

units and zero point. From equation (1), the logarithm of the ratio of the frequencies with which an entropy-constrained agent chooses two feasible actions is equal to the difference in implied payoffs to the actions scaled by the behavior temperature. When agents' behavior reflects entropy-constrained expected utility maximization, it is possible to recover quantitative information about payoffs from observations of behavior.

In scenarios where it makes sense to regard agents as identical, and their payoffs as interpersonally comparable, the cardinal properties of payoffs permit quantitative comparisons of payoff outcomes as parameters describing the institutional context of social interactions change. While conventional economic choice theory makes it possible, for example, to judge whether a given institutional structure will lead to a Pareto-efficient outcome which exhausts all possible increases in payoffs, the entropy-constrained approach supplies quantitative information about the degree to which any outcome approximates Pareto-efficiency. This quantitative information can be expressed either in the form of an estimate of the unrealized potential payoff gains in some particular outcome, or in the form of an estimate of the probability that two randomly chosen agents could make a mutually advantageous exchange.

Publisher's Note The EPJ Publishers remain neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

1. D. Kahneman, J.L. Knetsch, R.H. Thaler, *J. Econ. Perspect.* **5**, 193 (1991)
2. D.L. McFadden, *Ann. Econ. Soc. Meas.* **5**, 363 (1976)
3. J. Roberts, H. Sonnenschein, *Econometrica* **45**, 101 (1977)
4. R. Sethi, *J. Econ. Behav. Organ.* **29**, 113 (1996)
5. C.A. Sims, in *Rational inattention: a research agenda* (2012), pp. 1–22
6. A. Tversky, D. Kahneman, *Q. J. Econ.* **106**, 1039 (1991)
7. H.P. Young, *Econometrica* **61**, 57 (1993)