Slide 1 of 17

Lecture 2 Entropy-constrained

behavior

Economic applications

Entropy-constrained behavior

- Informational entropy and related concepts are central to statistical physics and related fields, but relatively exotic in economics.
- One recent strand of economics literature that begins to explore the possible implications of informational entropy for economic theories is the *entropy-constrained behavior* model. A version of this model appears in the literature

under the name *rational inattention*.

- This model applies information constraints to the theory of individual (or firm or household) behavior.
- The entropy-constrained behavior model can apply to Marxist concepts such as a capitalist's pursuit of profit, or to marginalist problems like consumer choice.

The entropy-constrained behavior model 1

- An “individual” (household, firm, person) decision is represented as maximizing an objective function, or *payoff*, $u[.]: X \to R$ defined over a domain of *actions, X*, which may be finite, countably infinite, or a continuum, possibly subject to some constraints.
- For simplicity we will subsume constraints such as technology or budget constraints into the action set X. Then the decision-maker is abstractly represented by the mathematical programming problem:

$$\underset{x \in X}{\text{Max}}\, u[x] \tag{1}$$

- This has the solution:

$$\hat{x}[X; u] = \text{ArgMax}[u[x] \mid x \in X] \tag{2}$$

Objective functions

- Practically speaking, the objective function represents something that theory suggests is predictable or repeatable in behavior.
- For example, Adam Smith argues that producers free to choose their occupation will move toward occupations they regard as more ‘advantageous”, typically by weighing the money income of an occupation against its disadvantages in terms of training, unpleasantness, and effort.
- Smith also argues that capitalists free to move

capital among different lines of production will seek the highest profit rate.

- Some political-economic philosophies such as utilitarianism regard objective functions as indexes of welfare and call them *utility functions*. It is possible to remain agnostic on this point.

Slide 5 of 17

The entropy-constrained behavior model 2

- We could also think of the individual as choosing a *mixed strategy* $f[.]: X \to R$ over the action set, and maximizing her *expected payoff.* In this case $f[.]: X \to R_{\geq 0}$, $\int f[x]\, dx = 1$ is a frequency distribution over actions:

$$\underset{\{f[x]\geq 0 \mid \int_X f[x]\, dx=1\}}{\text{Max}} \int f[x]\, u[x]\, dx \qquad (3)$$

- When X is a finite or countable set, we interpret the integral sign as a summation, and the frequency as a probability mass distribution.

- The solution to (3) is always to choose the utility-maximizing action:

$$\hat{f}[x; X, u] = \text{DiracDelta}\left[x - \hat{x}[X, u]\right] = \delta\left[x - \hat{x}[X, u]\right] \quad (4)$$

  - The *Dirac Delta* function $\delta[x]$ puts unit weight on 0 (thus on the utility-maximizing choice, $\hat{x}[X, u]$. ($\int \delta[x] \, dx = 1$, $\delta[x \neq 0] = 0$.)

Thermodynamics and behavior

- This type of theory raises the eyebrows of statistical physicists (and has since Boltzmann tried unsuccessfully to explain the reason to Walras), because the entropy of the Dirac Delta distribution is zero, due to the fact that it puts all the frequency weight on one point.
- In statistical physics zero entropy corresponds to a temperature of *absolute zero*, which physicists regard as an abstraction unattainable in the real world.

Thus economics seems to be operating from the start under extreme and unrealistic informational assumptions.

Slide 7 of 17

Entropy-constrained behavior

- One (not necessarily the only way) to introduce more realism (in the thermodynamic sense) into the economic behavior model would be to constrain the individual's mixed strategy frequency distribution over the action set to have a minimum informational entropy.
- If the individual is constrained to maximize expected payoff under a constraint on the entropy of the action distribution, her action distribution satisfies the following optimization, where $H_{\min}$ is the minimum entropy:

$$\underset{\left\{f[x]\geq 0 \,\middle|\, \int_X f[x]\,dx=1\right\}}{\text{Max}} \int f[x]\, u[x]\, dx$$

$$\text{subject to } H[f] = -\int f[x]\, \text{Log}[f[x]]\, dx \geq H_{\min}$$

- The objective function here is a linear, and hence concave function of the frequencies, and the constraint set is defined by the strictly concave entropy function, and is a convex set in frequency space.

Entropy-constrained behavior 1

- The constrained-entropy problem (5) is a well-behaved problem, with a Lagrangian (taking T to be the Lagrange multiplier or shadow price on the entropy constraint):

$$\mathcal{L}[f;\, T] = \int f[x]\, u[x]\, dx + T\left(-\int f[x]\, \text{Log}[f[x]]\, dx - H_{\min}\right) \quad (6)$$

  - The Lagrangian is the difference between the expected payoff and the entropy of the mixed strategy weighted by the Lagrange multiplier T, which we will call the *behavior temperature*.

- This difference is sometimes called the *free payoff*, or *free utility* by analogy with the concept of *free energy* in thermodynamics.

The entropy-constrained behavior problem 2

- The first-order conditions from the Lagrangian, (6), which are necessary and sufficient to characterize a global maximum over the behavioral frequencies, are:

$$u[x] - T\,(\text{Log}[f[x]] + 1) = 0 \tag{7}$$

- Thus the expected payoff-maximizing mixed strategy has $f[x] \propto e^{\frac{u[x]}{T}}$. The normalized frequency is the *Gibbs* or *logit quantal response* distribution:

$$f[x;T] = \frac{e^{\frac{u[x]}{T}}}{\int_X e^{\frac{u[x]}{T}} \, dx} \quad x \in X$$

- The Gibbs distribution describes any maximum entropy system constrained by moments of the frequency distribution of some measurable characteristic of a system.

Properties of the maximum entropy-constrained expected payoff

- The frequency with which the individual performs action *x* is proportional to the exponential of the payoff of *x*.
- The dispersion of individual behavior depends on the temperature-like variable *T*. At "absolute zero", where *T* –> 0, the individual's action becomes more and more concentrated on the payoff-maximizing choice.
- As long as *T* > 0, the individual will perform every available action *x* with some positive frequency.

Given $T > 0$ it would be possible to infer the utility of an action $u[x]$ from the frequency with which individuals were observed performing that action, up to an affine transformation.

- There is a *qualitative difference* in predicted behavior when behavior temperature is positive compared to payoff-maximization without an entropy constraint.

Cardinality of mixed-strategy frequencies

- An *affine transformation* of payoff consists of multiplying by a scalar $b > 0$, and adding a constant, a:

$$v[x] = a + b\, u[x] \tag{9}$$

- The maximing entropy-constrained mixed strategy frequencies for the transformed payoffs will be:

$$g[x; T, v] = \frac{e^{\frac{v[x]}{T}}}{\int e^{\frac{v[x]}{T}}\, dx} =$$

(10)

$$\frac{e^{\frac{a+b\,u[x]}{T}}}{\int e^{\frac{a+b\,u[x]}{T}}\,dx} = \frac{e^{\frac{a}{T}}\,e^{\frac{b\,u[x]}{T}}}{e^{\frac{a}{T}}\int e^{\frac{b\,u[x}{T}}\,dx} = f\left[x;\,\frac{T}{b},\;u\right]$$

- Thus the (observed) frequencies of an action determine payoffs up to a multiplicative factor, $\frac{T}{b}$, which can be regarded either as expressing the intensity of the individual’s preference over the actions, or the behavior temperature at which she operates.

lide 12 of 17

Frequency of actions and intensity of preference

- Taking logarithms of the logit frequency:

$$\text{Log}[f[x, T]] = \frac{u[x]}{T} + \text{Log}[Z[X, u, T]] \quad (11)$$

- Thus the difference of the logarithms of observed frequencies are differences of payoffs scaled by the decision temperature:

$$\text{Log}\left[\frac{f[x_1, T]}{f[x_0, T]}\right] = \frac{u[x_1] - u[x_0]}{T} \quad (12)$$

lide 13 of 17

Partition function

- The normalizing factor is the *partition function*:

$$Z[X, u, T] = \int_X e^{\frac{u[x]}{T}} dx \tag{13}$$

- The partition function for Gibbs distributions has some useful formal properties. For example:

$$\frac{\partial \text{Log}[Z[X, u, T]]}{\partial(1/T)} = \frac{\partial Z/\partial(1/T)}{Z} = \int_X u[x] \frac{e^{\frac{u[x]}{T}}}{Z} = \int_X u[x] f[x, T] = \overline{u} \tag{14}$$

The higher derivatives of the log partition function return the higher moments of the Gibbs distribution.

lide 14 of 17

Utilitarianism

- Utilitarian philosophers, including Jeremy Bentham and William Stanley Jevons, argued that individual behavior revealed individual welfare in the form of *utility*, and that welfare was comparable among individuals.
- The welfare of society could, according to utilitarian theory, be represented as the sum of individual welfares:

$$U^{\text{society}} = \sum_i u^i[x^i] \tag{15}$$

- Utilitarian ethics require maximization of this social welfare.
- With the common utilitarian assumption of diminishing marginal payoff to income, social welfare, in the absence of other constraints, is increased by equalizing incomes.

Interpersonal utility comparisons

- Any expected payoff theory implies the measurability of intensity of preference, and therefore has a partially cardinal character.
- If we adopt the utilitarian perspective that behavioral payoffs also represent welfare, then this theory suggests a partially cardinal measurement of welfare.
- In order for welfare to be *interpersonally comparable*, two further assumptions are necessary.

The first is that the individuals to be compared have the same intensity of preferences, that is the same b/T.

- The second is that we can assume that some particular state, x_0, gives the individuals the same welfare.

Risk of death

- One way to make the interpersonal comparison of welfare operational would be to assume that all people are equally well-off when they are dead.
- Since people frequently make choices that involve a risk of death, in principle observed behavior in these risky situations could establish an interpersonal welfare scale.

The dogma of non-interpersonal welfare comparisons

- Despite its roots in utilitarian thinking, marginalist economics shied away from the egalitarian implications of utilitarian philosophy.
- The mainstream marginalists insisted that interpersonal comparisons of welfare were impossible due to the inherently subjective aspects of welfare.
- The entropy-constrained behavior model verges on infringing the dogma of the non-comparability

of individual welfare.