EC 490

Lecture 18: Consolidation of the neuroeconomics of valuation, as of 2008

1. By a 'back-pocket' neuroeconomic model of choice, Glimcher means a general form of model from which each new investigation can begin, and which each investigator can aim to refine and elaborate upon. He makes clear that he expects a boot-strapping process to occur here; that is, expects that efforts to refine the model will ultimately lead to its replacement by a better one. He has, he says "certain knowledge that the details of the following framework are wrong." I suggest that in interpreting this remark we emphasize the word 'details'. I think there is a good possibility that the general form of the model will be vindicated as more evidence comes in. (In general, I'm unpersuaded that science teems with Kuhnian discontinuities.)

2. The most basic claim of the model is that choice occurs at the neural level through the chaining together of two distinct processes. First, a valuation mechanism ranks sets of options. Then a choice mechanism selects items from such sets and directs actions aimed at consuming the items in question. The basic evidence for this two-stage model is straightforward: repeatedly observed groups of neurons in VS and PFC fire in patterns that linearly co-vary with previously learned values of rewards even under conditions when animals face no relevant choices among actions. Distinct neural areas, especially in parietal cortex (at the top of the head), are active in choice, and receive input from the valuation system.

3. Glimcher uses economic theory to define objects that "serve … as mapping rules that connect existing theoretical tools to the empirical measurements of neuroscience." The central such relationship is the interpretation of value-in-the-brain as,

literally, utility (or something conceptually resembling it that often exactly mimics it – see below). Note that this was an empirical discovery, because it might not have been so. In organisms without brains, value is probably directly identifiable with maximization of some chemical magnitudes.

4. You are urged to carefully read what Glimcher says about utility theory on p. 507 and try to remember it when you hear nonsense about what 'neoclassical' economists supposedly believe. The truth and importance of the following remarks can't be over-emphasized: "these axioms are not some set of strange and arbitrary assumptions about how people *must* behave … The axioms are a statement not about people (or the brain) in any sense; the axioms are a precise definition of a theory." 'Theory' is here being used in the sense that mathematicians intend, and about which philosophers and others are often gravely confused. A theory is a device that precisely and unequivocally defines the relationships within a family of empirical models. It is not an elaborate conjecture about how the world might be.

5. Two key properties of utility are:

   (i) Utility is ordinal. This remains true even when we cardinally relate utilities by calibrating them on a common scale by using them to represent subjects' choices over lotteries with known probabilities of outcomes. Therefore, if value-in-the-brain is utility-like, then value-in-the-brain *cannot* be identified with a magnitude of any neural quantity, such as firing rates. Of course, value-in-the-brain could be a *function of* various neural magnitudes.

   (ii) Human choice is often inconsistent with the idea that a person maximizes a single utility function. This does not refute the idea that value-in-the-brain is utility-like. It merely shows us that the relationship between value-in-the-brain and behavior is not simple (which is hardly surprising).

Let us pass over the third property of utility listed by Glimcher, its identification with welfare. It is entirely implausible to identify value-in-the-brain with personal welfare (and Glimcher doesn't propose to do so). In my opinion, it is implausible to identify personal welfare with *any* property of an individual that is completely abstracted from her social environment. So we should never expect neuroeconomics to *directly* inform us about welfare. (I think that Glimcher agrees with this, for a different reason that might or might not be compatible with mine.)

6. Glimcher defines a new, distinctively neuroeconomic, object he calls *subjective value* (SV). SVs are real numbers ranging from 0 to 1000. Their units are action potentials per second. They are defined as the mean firing rates of specific populations of neurons, so they are linearly proportional to the BOLD signals measured in the populations in question. SVs are always stochastically consistent with choice, even when expected utilities are not. (Call this 'Property *'.) SVs have a unique reference-dependent anchoring point called the baseline firing rate. All SVs are encoded cardinally in firing rates relative to this baseline.

7. I am doubtful that any measurable magnitude has all of these properties *if the terms in which the properties are defined are given their standard, everyday interpretations*. In particular, I'm doubtful that any neural magnitude has property *, *unless choice is defined in neural terms.* Someone can so define choice if they like; but then a new term will be needed for what economists have traditionally meant by 'choice'. I'll come back to this issue at the end of the lecture.

8. Glimcher's reason for denying that neuroeconomics directly tells us about personal welfare is that SV doesn't respect the independence axiom of expected utility theory. That strikes me as a potentially deep and enlightening insight.

9. Note that SVs cannot take negative values, and that SV has finite range and variance. These properties will rule out what would otherwise be equilibria in various cases of neuroeconomic dynamics.

10. Glimcher proposes that "one central goal of neuroeconomics should be to develop a complete theory of SV." This is very helpful. If we provisionally accept this, then we can say that we know what is at stake when we wonder about the ultimate value of neuroeconomics.

11. Now some values are defined.

    (i) First, the *relative SV* of an option *j*:

    $$RSV_j = \frac{SV_j}{\sum SV_i + c}$$

    where *i* is the set of all options in the choice set, and *c* is an empirically measurable normalization constant derived from neuroscience. Glimcher argues that choices between options are made by comparison of their RSVs.

    (ii) Next, *obtained SV* (experSV): the SV of the state of the world that results from a choice. "The neural location of ExperSV is not known, though the activity of dopamine neurons provides overwhelming evidence that it is present as one of the midbrain inputs to those neurons." Sure: dopamine neurons couldn't implement TD learning otherwise.

    (iii) Reward prediction error, in terms of values defined so far:

    $RPE = \alpha(SV_{forecast} - ExperSV)$

12. Glimcher suggests that there are two sources of stochasticity relevant to choice: in valuation, corresponding to random utility distributions, and in the choice mechanism,

corresponding to the old game-theoretic notion of a trembling hand.

13. *SV variability* is a random term drawn from a Guassian distribution and added to mean SV to yield SV.

14. *Cortical noise before choice* is a random term drawn from a Poisson distribution and added to RSV as a stochastic time series before choice. It reflects the physics of cortical neurons.

15. Pulling all this together, the SV of an option *j* is the sum

$$SV_j = \frac{\sum_i \omega_i x_{ij}}{\sum_i \omega_i}$$

where *i* indexes each neuron in the brain, x_i is the firing rate of the *i*th neuron, and ω_i is a weight ranging from 0 to 1 describing the additive contribution of the *i*th neuron to the SV of *j*.

16. Glimcher identifies two questions that he calls "paramount" for a scientist looking for some SV_j in the brain:

    (i) Can we identify a firing rate (or a proxy for it, e.g. BOLD activation) of some neurons that is linearly correlated with the utility of actions or options that the subject chooses – in those cases when utility predicts choice?

    (ii) What is the smallest group of neurons that can maintain this linear correlation (i.e., the smallest group for which $\omega_i \neq 0$.

    (What are we to do when utility does *not* predict choice?)

17. Glimcher then ventures a strong hypothesis: all the neurons we need to study to find the SV for any object are in VS and PFC.

18. Glimcher hypothesizes that dopamine neurons directly compute SV under conditions indicated by the following model:

$$SV_{jt} = SV_{j(t-1)} + \alpha(ExperSV - SV_{j(t-1)})$$

19. Aspects of the computation of reinforcement effects occur in basal ganglia. However, Glimcher interprets the weight of fMRI and other evidence as indicating that the SV of an action or good is encoded by the mean activity in the medial PFC and the VS. He speculates that medial PFC stores goods valuation and VS stores action valuation.

20. Because neurons have finite rates of stochasticity, SV estimates must be drawn from an underlying distribution. Therefore, their estimation is said to be more closely analogous to random utility estimation than to von Neumann – Morgenstern calculation. However, random utility models share the von Neumann – Morgenstern assumption that the decision-maker has perfect discrimination capability. Random utility models are motivated by the notion that the *analyst* may have incomplete information. Manski (1997) identifies four different sources of analyst's uncertainty: unobserved alternative attributes, unobserved individual attributes (called 'unobserved taste variations'' by Manski), measurement errors in proxy, or instrumental, variables.

    Utility is modeled as a random variable in order to reflect this uncertainty. More specifically, the utility that individual *i* associates with alternative *a* is given by

$$U_\alpha^i = V_\alpha^i + \varepsilon_\alpha^i$$

    where V_α^i is the deterministic component, and ε_α^i is stochastic (as far as the analyst is concerned). The probability that alternative α is chosen by decision-maker *i* within choice set

C is

$$P_C^i(\alpha) = P[U_\alpha^i = U_\alpha^i \max_{b \in C} U_\alpha^i$$

Presumably Glimcher assumes that the uncertainty in SV estimation lies in the world rather than in the restricted knowledge of the modeler. But this isn't inconsistent with his borrowing random utility theory. Just think of the brain, trying to match its responses to the actual distribution of rewards, as the analyst.

21. What about the choice system? Glimcher describes the substantial discoveries of his lab on the selection of eye movements, where specific movements are associated with different rewards, in monkeys. According to him, these depend on use of the SV signal encoded in VS by parietal (area LIP) neurons. He concedes, however, that at this point relatively little is known about the computation of choice outside of this restricted context.

22. Glimcher argues that evidence bears against 'multiple self' models. By this he refers to models that partition 'emotional' or 'limbic' decision-making from 'cognitive' or 'frontal' decision-making. The weight of evidence suggests that both systems traditionally regarded as emotional, and systems traditionally regarded as cognitive, jointly contribute to the computation and subsequent use by various brain areas of single SV signals. Note that this is not an argument against a *different kind* of multiple-self model as promoted by, e.g., Ainslie. Such 'picoeconomic' models distinguish between personal (or 'molar') utility, which may be partly calculated by the person in conjunction with external systems, and value-in-the-brain. A picoeconomist can and should grant that SV or something much like it is a crucial *input* to personal choice, but will also allow for other sorts of input,

many of which are simply 'choice-governing tracks' laid out in the subject's cultural / social / market environment.

23. I think Glimcher will be suspicious of picoeconomic models, for reasons unrelated to his grounds for skepticism about dual-self models. There are deep philosophical issues here, related to the question of whether 'choice' must denote an *inboard* process that describes an agent in isolation. There is no doubt that the *internalist* about choice takes the view that everyday intuition finds more natural. As a matter of methodology, I also think that this is the view that neuroeconomists should pursue as far as they can. After all, the *brain* can only play a role in choice insofar as something like SV representation is involved.

24. But, yes, the externalist is committed to the strange idea that there can be elements of 'choice' that don't involve the brain at all. See Clark, *Being There*, 1997; Ross, *Economic Theory and Cognitive Science: Microexplanation*, 2005. Externalists have two bodies of reasoning on their side: recent work in cognitive science and robotics that emphasizes the crucial role of *unrepresented* external structure in constraining behavior, and the tradition in economics that identifies 'rational choice' with *what can be rationalized* rather than with *what is explicitly calculated*. It is the force of this tradition that leads me to resist the suggestion that we define choice in neural terms (see bullet 7 above). But the best means of testing this suggestion is exactly Glimcher's program: try to provide neural foundations for as much of microeconomics as possible, and see what if anything is left over at the conclusion.

THE END