Distracted Learning

(by Axel Anderson and Andrea Wilson)

We introduce a model of distracted learning mitigated by an optimal N-state memory buffer. Unlike the finite automata literature (eg Wilson, 2014), the model is founded on a fully Bayesian decision maker (Dug), who continuously updates his beliefs given an Ito signal process. When the termination shock hits, he must choose a binary action, earning a higher payoff if his action matches the binary state. The novel difference is that Dug can get distracted. When a distraction shock occurs, he can only recall which one of N intervals (memory states) contained his belief when the distraction shock hit. He maximizes his expected payoff over interval partitions $\{p_n\}$ of [0, 1], and his post-distraction beliefs $\{q_n\}$. For instance, with N=2 memory states, the model is fully described by a triple (p_1,q_1,q_2).

Dug is optimally indifferent between memory states at the threshold between memory states. His unoptimized value obeys smooth pasting at the $\{p_n\}$, but optimal values are "super-smooth pasted" (matched second derivative). Dug is always harmed by distraction shocks, and gains at a termination shock if and only if he is sufficiently certain of the state – in other words, his value $V^*(p)$ falls below the myopic value at extremes:

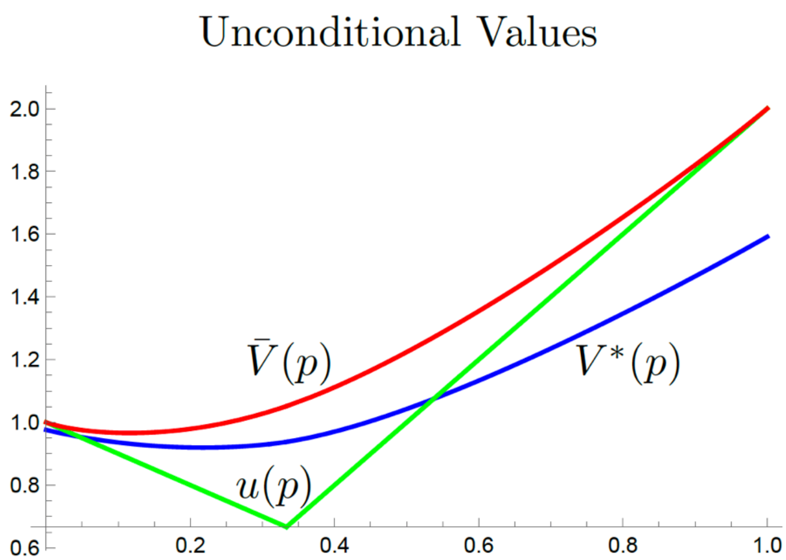

His initial value rises in the number of memory states and falls in the rate of termination shocks. Two memory states are sufficient to secure the full information value as termination shocks become vanishingly rare.

We consider two extensions: allowing the DM to choose the optimal number of memory slots (demand for memory) or the optimal precision of the Ito process (demand for information) given some increasing cost function.