Notes on "The Optimal Level of Experimentation"

(Giuseppe Moscarini and Lones Smith, 2001)

- Sam gets one-shot payoffs π_a^θ from action $a \in \{A, B\}$ in state $\theta \in \{0, 1\}$. Thus, the expected payoffs from action a are $\pi_a(p) = p\pi_a^1 + (1-p)\pi_a^0$.

- Action A is best in state 0 and action B is best in state 1. So Sam's one-shot payoff is $\bar{\pi}(p) = \max(\pi_A^0(p), \pi_B^1(p))$, with indifference at $p = \hat{p}$.

- Sam has an initial experimentation phase: He sees in time interval $[0, t)$ a sample path realization of an Ito process. Chernoff (1972), "Sequential Analysis and Optimal Design" allowed a continuous time observation process (X_t), where $dX_t = \theta dt + \sigma dW_t$, for an Ito process (W_t).

- Show: Bayes rule asserts $dP_t = \alpha P_t(1-P_t)d\bar{W}_t$, for an Ito process $\bar{W}_t$.

- If experimentation has flow cost $c > 0$, the optimal solution is a stopping time T: continue until the posterior p_T exits $[\underline{p}, \bar{p}]$, where $\underline{p} < \hat{p} < \bar{p}$.

- As posterior belief variance with n conditionally iid signals is scaled by n, call the level of experimentation n if $dP_t = \alpha\sqrt{n}P_t(1-P_t)d\bar{W}_t$.

- Assume Sam is impatient and experiments at variable intensity. With a constant marginal cost of experimentation, Sam intuitively stops in an "instant", when P_t hits $\underline{p}$ or $\bar{p}$. So assume cost $c(n)$, with $c', c'' > 0$.

$$V(p_0) = \max_{T, n_t} E\left[\int_0^T -c(n_t)e^{-rt}dt + e^{-rT}\bar{\pi}(P_T)\right] \text{ s.t. } dP_t = \alpha\sqrt{n}P_t(1-P_t)d\bar{W}_t$$

- Optimal stopping needs value matching and smooth pasting (Chernoff).

- Since $E[dP_t] = 0$, the Bellman equation for optimal control reduces to

$$rV(p) = \max_{n \geq 0} -c(n_t) + \tfrac{1}{2}np^2(1-p)^2V''(p) \Rightarrow \text{FOC } c'(n) = \tfrac{1}{2}p^2(1-p)^2V''(p)$$

- Combining, $rV(p) = -c(n) + nc'(n) = g(n)$, where $g'(n) = nc''(n) > 0$. So Sam's return rV is the producer surplus of experimentation $g(n)$.

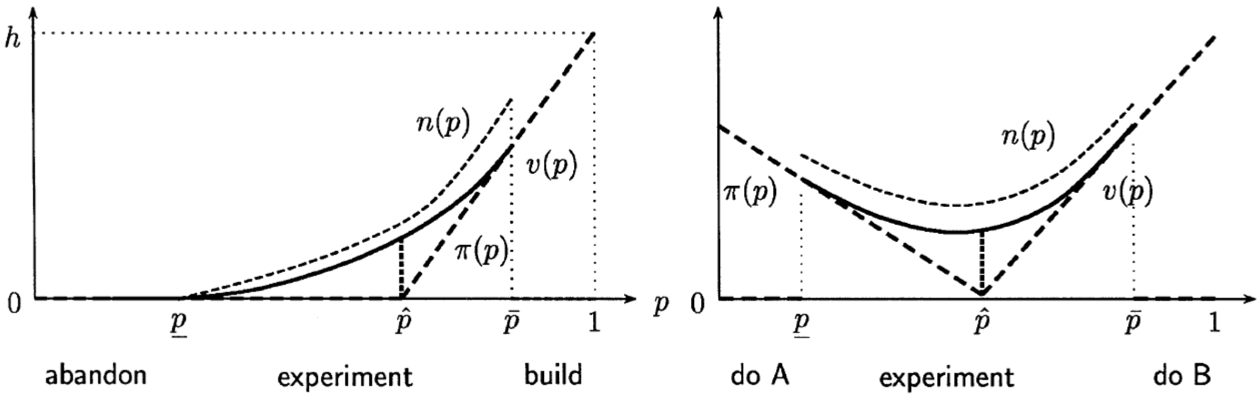

- The optimal level of experimentation is the inverse producer surplus $g^{-1}(rV(p))$, and so shares the ordinal shape of the value. In the classic case of an affirmative build action if $\theta = 1$, and an abandon action otherwise, the level is increasing in p. In the general two action case A or B with a $\vee$-shaped terminal payoff frontier, the level is U-shaped.

- Example: With quadratic costs, $c(n) = n^2$, we have $g(n) = n^2$, and thus $n(p) = \sqrt{V(p)}$.

- Loose applications: Our efforts to finish any project ramp up near the end, to optimally back-end the costs, and diminish their present value.

- The paper develops testable predictions of the model. Notably, the optimal experimentation level increases in the interest rate r near the stopping thresholds $\underline{p}$ and $\bar{p}$ (Proposition 5(e)). Namely, the demand for experimentation in a dynamic world is "backward-bending" as a function of the interest rate near the end, as Sam rushes to finish sooner.[1]

- Years later, Raj Chetty (2007)[2] unwittingly published a two-period version of Moscarini and Smith, albeit with no control of precision, just a stopping problem (today or tomorrow). His paper's main result was the backward-bending demand for investment when the investment payoff is uncertain, and the cost is sunk once undertaken.

[1] Also, the stopping thresholds $\underline{p}$ and $\bar{p}$ shift in (so one stops earlier, as is natural). The Proposition has a typo, saying they shift out. The appendix has the correct statement and proof (Claim 2).

[2] *Review of Economic Studies*, "Interest Rates, Irreversibility, and Backward-Bending Investment", Volume 74, Issue 1, pages 67–91. It did not cite Moscarini and Smith.