Sequential Sampling Models

Mark Dean

Behavioral Economics G6943
Fall 2016

Sequential Sampling Models

- In the last couple of lectures we have considered models in which the form of information acquisition has been very flexible
- Today we are going to consider a set of models in which the process is much more constrained
  - Sequential sampling models
- These models have been hugely influential in the psychology literature
- Increasingly so in the economics literature
  - On the first day of this year's Cowles' theory conference, 5 of the 6 papers were about sequential sampling models (!)
  - Including papers by Yeon Koo Che and Mike Woodford

The Basic Idea

- Consider choosing between two alternatives
  - Typically a perceptual task - e.g. Left or Right in a dot motion task

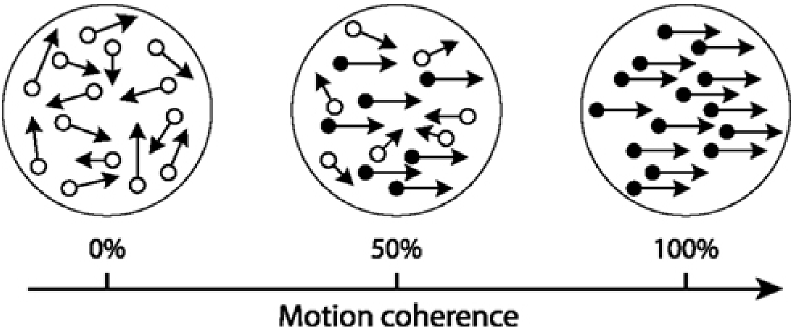
0%
50%
100%
Motion coherence

The Basic Idea

- Consider choosing between two alternatives
  - Typically a perceptual task - e.g. Left or Right in a dot motion task
  - More recently applied to value based decision making - e.g. Apple or Orange
- Over time, evidence accumulates about each alternative
  - Observe dots moving left or right
- This evidence is noisy
- The DM must construct a rule that tells them when to stop gathering more information and make a choice
- This is (basically) the class of sequential sampling models (SSMs)

Why So Popular

- There are (I think) four reasons that these models have proved so popular
  1. Intuitive plausibility
  2. Biological plausibility
  3. Links to optimality
  4. Ability to predict relationship between choice and reaction time

Sequential Sampling Models

- There is a huge variety of sequential sampling models
  - Relative vs absolute stopping rules
  - Discrete vs continuous accumulation of evidence
  - Fixed vs collapsing bounds
- See Ratcliff and Smith [2004] for a taxonomy
- We will first consider the (Drift) Diffusion version of the model
  - See Shadlen et al. [2007]

Drift Diffusion Model

- An SSM has three components
  1. The process by which evidence is accumulated
  2. The bounds that govern the decision
  3. Reaction (i.e. non-decision) time

- Evidence Accumulation
  - We will assume that in each period the DM receives a signal X
    - Distributed iid according to a distribution with mean μ and variance σ^2
  - Evidence is therefore of the form of a sequence $\{X_1, ..., X_n, ..\}$
  - The 'sum' of evidence is therefore given by a sequence $\{Y_1, ..., Y_n, ..\}$ where

$$Y_n = \sum_{x=1}^{n} X_n$$

  - Easy to move to continuous setting if convenient

$$\frac{dY}{dt} = \mu + N(0, \sigma\sqrt{dt})$$

- Boundary
  - We will assume that evidence accumulates until either $Y > A$ or $Y < -A$
- Reaction time
  - We will assume that, on each trial, there is a time period t_{nd} before which evidence starts to accumulate
  - This is drawn from a distribution with mean $\bar{t}_{nd}$
  - Basically a kludge to better fit the data

Drift Diffusion

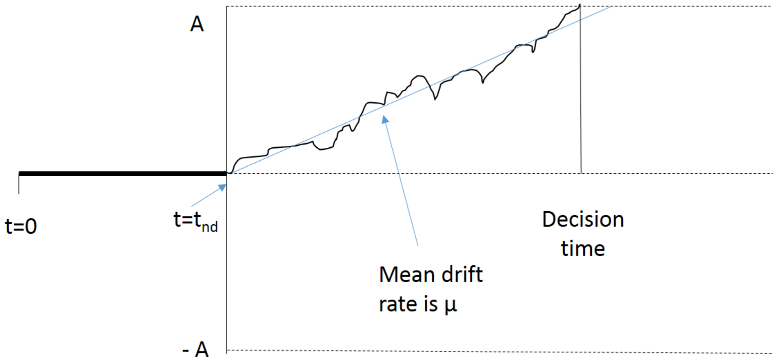

Drift Diffusion - Links to Experiment

- How do the parameters of the DDM link to the parameters of the experiment?
- Let C be the strength of the signal on a particular trial
  - e.g. the fraction of dots moving to the left
- μ is assumed to be an increasing function of C
- What about boundaries A?
  - Cannot depend on the C on a particular trial
  - But may depend on the expected distribution of C
  - May also depend on the costs of different type of error
  - See discussion of optimality
- Note, σ, μ and A are not separately identified
  - Will set $\sigma = 1$

Drift Diffusion - Predictions

- Let's assume for the moment that A is fixed
- What are the predictions of the model?
- Note that the data of interest is
  - Joint distribution of choices and reaction times
  - Conditional on the signal strength C
- We will go through this quickly
  - For details see Shadlen et al [2007]

Choice Probabilities

- First, we want to understand the probability of choosing each option as a function of C
- This boils down to calculating the probability that Y_t will first hit A or $-A$ as a function of the distribution of X
- A handy reminder: The moment generating function (MGF) of a random variable:

$$M_X(\theta) = E(e^{\theta X}) = \int f(x) e^{\theta x} dx$$

- Recall that the nth moment of the distribution is given by

$$\frac{d^n M_X(\theta)}{d\theta^n}\Big|_{\theta=0}$$

- Typically, a MGF will have two values of θ such that $M_X(\theta) = 1$: 0 and θ_1

- **Step 1:** Note that the MGF of the unconstrained process Y_n is

$$M_{Y_n}(\theta) = M_X^n(\theta)$$

- **Step 2:** Let $\bar{Y}$ be the random variable which is the termination value of Y. Then

$$M_{\bar{Y}}(\theta) = P_+ e^{\theta A} + (1 - P_+) e^{-\theta A}$$

where P_+ is the probability of terminating at the the top boundary (this is what we want to find)

- **Step 3:** Define Wald's Martingale

$$Z_n = M_X^{-n}(\theta) e^{\theta Y_n}$$

- **Step 4:** Note that

$$E[Z_{n+1} | Z_n] = Z_n$$

i.e. Wald's Martingale is a martingale

- **Step 5:** Note that

$$E[Z_n] = E[M_X^{-n}(\theta)e^{\theta Y_n}] = M_X^{-n}(\theta)E[e^{\theta Y_n}] = 1$$

- **Step 6**: Define $\bar{Z}$, the 'stopped' version of Z

$$\bar{Z} = M_X^{-\bar{n}}(\theta)e^{\theta \bar{Y}}$$

  where $\bar{n}$ is a random variable

- **Step 7**: Apply the Optional Stopping theorem for martingales which states

$$E[\bar{Z}] = E[Z_n]$$

  and so

$$E[\bar{Z}] = E\left[M_X^{-\bar{n}}(\theta)e^{\theta \bar{Y}}\right] = 1$$

- **Step 8:** Recall that there is a value θ_1 such that $M_X(\theta_1) = 1$, and so

$$E\left[e^{\theta_1 \bar{Y}}\right] = 1$$

  - This is true as long as $E(X) \neq 0$ and X can take positive and negative

- **Step 9**: Note that this is the MGF for $\bar{Y}$ and so

$$\begin{aligned} P_+ e^{\theta_1 A} + (1 - P_+) e^{-\theta_1 A} &= 1 \\ &\Rightarrow \quad P_+ = \frac{1}{1 + e^{\theta_1 A}} \end{aligned}$$

- **Step 10:** Note that if we make some distributional assumptions about X we can solve for θ_1
  - e.g. if $X \sim N(\mu, 1)$ then

$$\theta_1 = -2\mu$$

Choice Probabilities

- If we then assume that μ is linearly related to stimulus strength by

$$\mu = kC$$

  then this gives us

$$P_+ = \frac{1}{1 + e^{-2kCA}}$$

- This is basically the Logit choice function

- What about the distribution of the length of time until choice is made?
- Take derivatives of the following with respect to θ

$$E[M_X^{-n}(\theta)e^{\theta Y_n}] = 1$$

giving

$$E\left[e^{\theta\bar{Y}}\bar{Y}M_X^{-\bar{n}}(\theta) - e^{\theta\bar{Y}}M_X^{-1-n}(\theta)M_X'(\theta)\right] = 0$$

- Evaluate at $\theta = 0$, and recall that $M_X'(0) = \mu$ and $M_X(0) = 1$ gives

$$\begin{aligned} E[\bar{Y} - \bar{n}\mu] &= 0 \\ &\Rightarrow \quad E(\bar{n}) = \frac{E[\bar{Y}]}{\mu} \end{aligned}$$

- Note that

$$E[\bar{Y}] = P_+A + (1 - P_+)(-A)$$

- Subbing in for P_+ and applying some magic gives

$$E[\bar{n}] = \frac{A}{\mu} \tanh\left(\frac{-\theta_1 A}{2}\right)$$

- Which, under normality, becomes

$$E[\bar{n}] = \frac{A}{\mu} \tanh\left(\mu A\right)$$

or

$$E[\bar{n}] = \frac{A}{kC} \tanh\left(kCA\right)$$

- This an expression for the expected number of 'steps' before a choice is made

Speed/Accuracy Trade Off

- So far we have derived expressions for the marginal distribution of choice accuracy and reaction time
- However, the model makes predictions about their *joint* distribution
- This is Speed/Accuracy trade off
  - Are ‘correct’ choices quicker or slower than incorrect ones?
- We have to be careful here about exactly what we mean
  - For a fixed difficulty: No parameters of the model change
  - For unanticipated changes in difficulty: C changes, meaning μ changes
  - For anticipated changes in difficulty: C changes, μ changes, and A may also change

Speed/Accuracy Trade Off

- Case 1: Fixed difficulty
- Surprisingly there is **no** speed accuracy trade off in the case in which
  - Bounds are equidistant from the starting point
  - Momentary evidence accumulation is normal
- For every path that goes to the upper bound there is an equivalent path that goes to the lower bound
- The mean reaction time is the same for correct and incorrect responses

- Case 2: Unanticipated changes in difficulty
- An increase in signal strength will
  1. Increase the probability of a correct choice

$$P_+ = \frac{1}{1 + e^{-2kCA}}$$

  2. Decrease reaction time

$$\begin{aligned} E[\bar{n}] &= \frac{A}{kC} \tanh(kCA) \\ \Rightarrow \quad \frac{dE[\bar{n}]}{dC} &= -\frac{A}{kC^2} \tanh(kCA) + \frac{A^2}{C} (1 - \tanh^2(kCA)) < 0 \end{aligned}$$

- Speed/Accuracy relationship **positive on average**: Faster choices are more accurate
- However, controlling for difficulty there will be no relationship between speed and accuracy

Speed/Accuracy Trade Off

- Case 3: Anticipated changes in difficulty
- Change in accuracy and reaction time will depend on change in μ and change in A
- Requires a model of where A comes from.
- See next section....

Optimality

- So far, we have described a procedure for decision making
- We have no idea whether it is in fact any good
  - i.e. is this a model of bounded rationality?
- The answer is yes, in the sense that there are problems for which this class of behavior is optimal.

Optimality and the Sequential Likelihood Ratio Test

- Consider the following problem
  - There are two states of the world, ω_1 and ω_2
  - In each time period you observe a signal γ the distribution of which is $f(\gamma|\omega_i)$
  - You have to identify the true state with an imposed level of accuracy
  - What decision rule minimizes the average number of observed signals
- Optimal solution consists of boundaries k_1, k_2 on

$$\frac{f(\gamma_1|\omega_1)f(\gamma_2|\omega_1)f(\gamma_3|\omega_1)....f(\gamma_n|\omega_1)}{f(\gamma_1|\omega_2)f(\gamma_2|\omega_2)f(\gamma_3|\omega_2)....f(\gamma_n|\omega_2)}$$

  Such that evidence is accumulated until the likelihood ratio goes above k_1 or below k_2
- This result dates back to Wald and Wolfowitz [1947]

Optimality and the Sequential Likelihood Ratio Test

- The sequential likelihood ratio test can be implemented as a diffusion model
- Take logs of the likelihood ratio

$$
\begin{aligned}
& \log\left[\frac{f(\gamma_1|\omega_1)f(\gamma_2|\omega_1)f(\gamma_3|\omega_1)....f(\gamma_n|\omega_1)}{f(\gamma_1|\omega_2)f(\gamma_2|\omega_2)f(\gamma_3|\omega_2)....f(\gamma_n|\omega_2)}\right] \\
= & \left[\log f(\gamma_1|\omega_1) - \log f(\gamma_1|\omega_2)\right] \\
& + \left[\log f(\gamma_2|\omega_1) - \log f(\gamma_2|\omega_2)\right] \\
& + \left[\log f(\gamma_3|\omega_1) - \log f(\gamma_3|\omega_2)\right] ..
\end{aligned}
$$

- So, defining $X_i = \log f(\gamma_1|\omega_1) - \log f(\gamma_1|\omega_2)$, the optimal stopping rule is to wait until $\sum_i X_i$ goes above $\log k_1$ or below $\log k_2$
- If γ is distributed log normally, the X_i will be distributed normally

- This can be mapped into a consumer choice problem
  - Choosing between two goods
  - The value of one good is θ_l the value of the other is θ_r
  - Evidence accumulates with a drift rate which is linearly related to $(\theta_l - \theta_r)$
  - Per signal cost c
- This fits into the above framework if there are only two possible 'states'
  - Either left is good and right is bad
  - Or visa versa
  - Difference in utilities is known
- See Fudenberg et al. [2016]

The Problem with Fixed Boundaries

- The model so far predicts either
  - No speed accuracy trade off at all if difficulty doesn't change
  - No speed accuracy conditional on difficulty if there are unanticipated changes in difficulty
- As we shall see this does not fit with the evidence
- Result comes directly from the fact that boundaries are fixed over time
- This in turn is a result of the assumption that there are only two possible states of the world
  - Intuitively, can never learn that two alternatives are hard to distinguish
  - i.e. learn that you are close to indifferent between two goods

Learning Indifference

- What if, instead, your prior is that the value of each alternative is drawn from some distribution
- What does optimal policy look like?
- Now, if you have not hit a boundary after a long time, it tells you that the drift rate is likely to be low
  - Implies difference in values is low
  - Value of further learning likely to be small
- This case is studied by Fudenberg et al [2016] and Tajima et al [2016]
- Turns out optimal policy is to have bounds that collapse over time

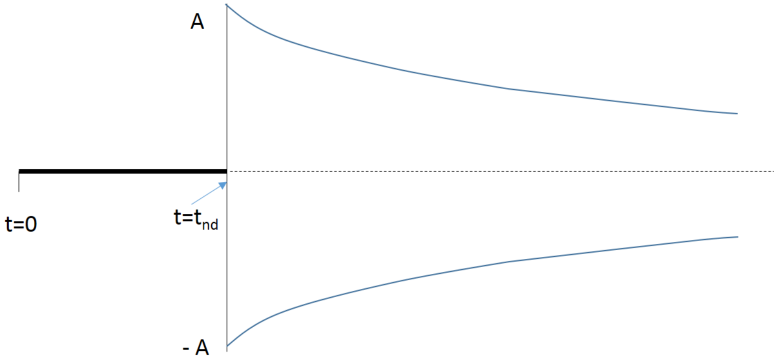
A
t=0
t=t_nd
- A

- This will lead to a positive correlation between speed and accuracy
  - On average across all decision problems
  - Conditional on difficulty
- For a fixed drift rate, hitting the boundary later increases the probability of error

Negative Speed Accuracy Trade Off

- Does the above mean that a positive speed accuracy trade off is inevitable
- No!
- Fix a difficulty level, and increase the rewards for making the correct decision
  - e.g. the experiments we saw in rational Inattention
- This will have the effect of increasing the boundaries A, while leaving C unchanged

$$P_{+} = \frac{1}{1 + e^{-2kCA}}$$

$$E[\bar{n}] = \frac{A}{kC} \tanh(kCA)$$

- Will increase both accuracy and reaction times
- So averaging over changes in reward level we will see a negative correlation between speed and accuracy

- The literature testing DDM type models is vast...
- ...and frankly I do not know half of it
- Good recent reviews include
  - Radcliff and Smith [2004]
  - Bogacz et al [2006] (also covers a lot of the theory well)
  - Radcliff and McKoon [2008]
- I will
  - Report some of the stylized facts from the perceptual literature
  - Discuss an application to economic decision making

Stylized Facts - Radcliff and Mckoon [2008]

1. Reaction time distribution is positively skewed
2. Increase in difficulty increases reaction time and decreases accuracy
3. Increase in difficulty increases positive skew
4. Response times for errors are often slower than for correct responses, even controlling for difficulty
    - But this can flip when accuracy is high or speed is emphasized
5. Emphasizing speed rather than accuracy reduces reaction time and increases error

Stylized Facts - Radcliff and Mckoon [2008]

- Can these facts be matched by the DDM with non-collapsing boundaries?
- Yes, if one allows for variability in drift rates and starting points
  - Importantly, not completely explained by apparent task difficulty
- This is sometimes called the 'full' DDM with 7 parameters
  - Mean and SD of the drift rate
  - Boundary
  - Mean and variance of the starting point
  - Mean and variance of non-decision time

- This comes directly from variance in the drift rate

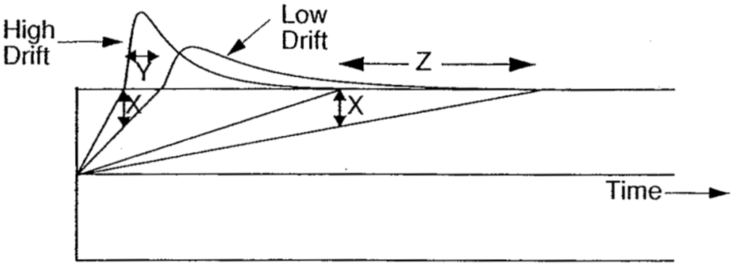

Effect of Difficulty

- Increase in difficulty decreases drift rate
- We showed that this decreases accuracy and increases reaction time
- It will also increases the skewness (by the argument on previous slide)

Speed/Accuracy Trade off

- As we discussed, if there is no variance in drift rate, model predicts no speed accuracy trade off
- However, if there is variance in the drift rate, this can make error trials slower than correct trials
  - This is an alternative to collapsing boundaries
  - Model fits tend to favor this approach
- How can the model capture the reverse effect?
- Variance in starting point

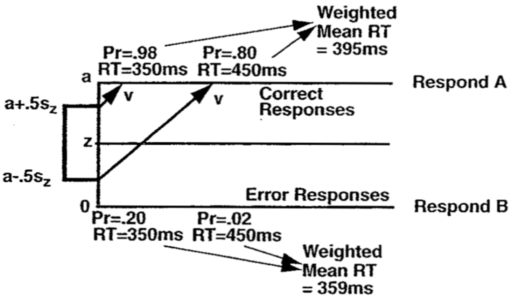

- As a proportion, higher fraction of errors come from starting points near that boundary
- Means they have lower response time on average

Effect of Emphasizing Speed vs Accuracy

- Emphasizing speed rather than accuracy is assumed to increase value of time relative to success
- Optimal response is to bring boundaries in
- Would lead to a reduction in accuracy and increase in response
- As seen in the data

DDM in Economic Choice

- Neuroeconomists have been very keen in using the DDM to fit economic choice
  - Particularly Rangel Lab
- Milosavljecic et al [2010] paragdimatic example
  - Subjects asked to rank 50 food items on a 5 point scale
  - Used to measure 'utility'
  - Then make 750 binary choices between randomly selected pairs
  - High and low time pressure conditions
  - One choice actualized at the end of the experiment

Table 1: Individual performance by condition, averaged over all values of *d*.

Subject	N	Accuracy (%)	Mean RT (S.E.M.) All Trials	Mean RT (S.E.M.) Correct Trials	Mean RT (S.E.M.) Error Trials
LOW TIME PRESSURE					
1	749	75.8	436 (4.77)	444 (5.73)	411 (7.83)
2	749	81.0	514 (4.03)	510 (4.16)	530 (11.46)
3	750	84.4	623 (5.44)	608 (5.77)	704 (13.37)
4	738	90.0	533 (5.89)	534 (6.02)	521 (22.78)
5	719	97.8	811 (10.26)	807 (10.22)	982 (98.29)
6	750	85.3	681 (6.11)	671 (6.34)	737 (18.60)
7	746	85.7	480 (6.23)	484 (6.37)	452 (17.53)
8	738	66.8	520 (6.73)	510 (6.67)	540 (13.10)
MEAN	742	83.2	574	578	552
HIGH TIME PRESSURE					
1	749	76.5	343 (2.88)	344 (3.24)	339 (6.26)
2	744	73.1	497 (3.13)	496 (3.64)	498 (6.14)
3	747	75.2	479 (4.13)	475 (4.40)	492 (9.94)
4	747	85.3	404 (3.64)	398 (3.76)	439 (11.26)
5	745	80.1	406 (2.44)	405 (2.57)	409 (6.57)
6	738	71.2	520 (3.70)	521 (4.45)	518 (6.60)
7	747	63.5	325 (2.61)	335 (3.35)	307 (3.92)
8	748	80.7	436 (3.26)	434 (3.49)	446 (8.49)
MEAN	745	75.7	426	426	426

DDM in Economic Choice

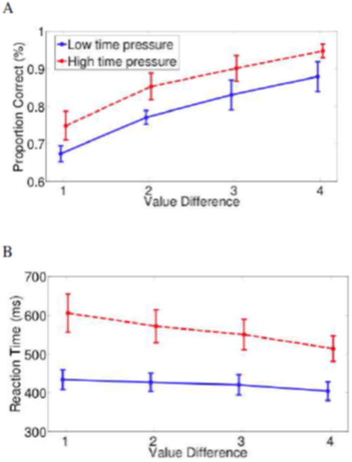

Summary

- Full DDM can do a reasonable job of fitting the data in both
  - Perceptual tasks
  - Choice tasks
- However, it needs the additional degrees of freedom provided by randomness in the
  - Drift rate
  - Starting point
- With these parameters, adding decaying boundaries does not improve fit
  - Interesting question: how to differentiate between stochastic drift rates and collapsing boundaries

Biological Plausibility

- One of the reasons that SSMs are so popular is that they seem to relate to actual neurological processes
- See for example
  - Gold and Shadlen [2007]
  - Bogacz [2007]
- Here is some evidence from Shadlen et al. [2007]
- Recording from various brain areas in monkeys during a dot motion task

Biological Plausibility

- Momentary accumulation of evidence is encoded in an area MT/V5
- Known from lesion studies and stimulation that this area is involved with eye movements
- Moreover, activity seems approximately linear in coherence
- Parameter estimates from neural data similar to those from behavior

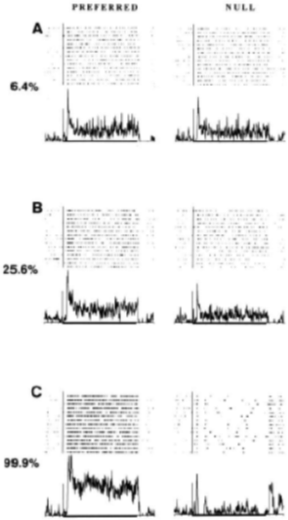
PREFERRED
NULL
A
6.4%
B
25.6%
C
99.9%

- A second area known as LIP appears to record accumulated evidence

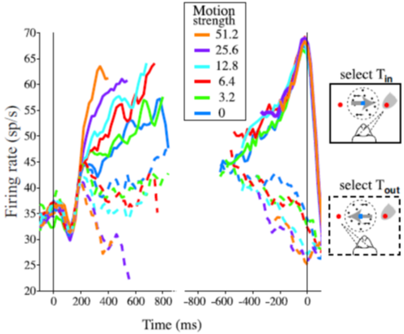

Summary

- SSMs provide a model that allows for joint predictions of reaction times and choice probabilities
- The full DDM provides a parsimonious way of modelling both perceptual and economic decisions
  - But potentially has a whiff of 'kludge'
- Other interesting extensions
  - Multiple options
  - The role of attention
  - Revealed indifference
- Interesting experimental avenue: chase down the implications of the optimal model.