Sequential Evidence Accumulation

Tomasz Strzalecki

Summer School in Economic Theory, IIAS Jerusalem, 2023

Response Times

- So far, given a menu we recorded *what* the agent chose
- Now: we also record *how long* the agent spends choosing
- In each instant the agent decides whether to stop and make a choice or delay the decision. For menu $A = \{\ell, r\}$ the decision problem is:

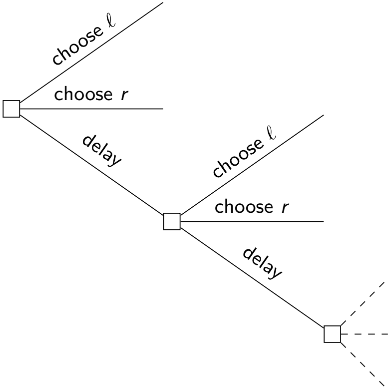

Benefits and Costs of Delaying Decisions

- Benefits: get more information
  - from outside: informative signals
  - from within: introspection/memory
- Costs:
  - opportunity cost of time
  - delaying consumption

Two Effects

Informational Effect (a.k.a. Speed-Accuracy Tradeoff):

- More time $\Rightarrow$ more information $\Rightarrow$ better decisions
  - seeing more signals leads to more informed choices
  - if we forced agent to stop at time t, make better choices for higher t
    - $\rightsquigarrow$ increasing accuracy

Two Effects

Informational Effect (a.k.a. Speed-Accuracy Tradeoff):

- More time $\Rightarrow$ more information $\Rightarrow$ better decisions
  - seeing more signals leads to more informed choices
  - if we forced agent to stop at time t, make better choices for higher t

    $\rightsquigarrow$ increasing accuracy

Selection Effect:

- Time is costly, so your decision to delay depends on how much you expect to learn (option value of waiting)
  - want to stop early if get an informative signal $\rightsquigarrow$ good decisions
  - want to delay if get a noisy signal $\rightsquigarrow$ presumably worse decisions

- Creates dynamic selection and can reverse the informational effect
  - if allowed agent choose t, make worse choices for higher t

    $\rightsquigarrow$ decreasing accuracy

Decreasing accuracy

The two effects push in opposite directions. Which one wins?

Stylized fact: Decreasing accuracy: if we group the universe of all decisions by the (endogenous) response time, then fast decisions are "better" and slow decisions are "worse"

- Well established in perceptual tasks, where "better" is objective
- Also in experiments where subjects choose between consumption items

Comment: The opposite is true in choice problems especially engineered by psychologists to contain trick questions where your first instinct is wrong (e.g. cognitive reflection test)

Observables

- S is the state space
- Time is discrete $\mathcal{T} = \{0, 1, 2, \dots\}$ or continuous $\mathcal{T} = [0, \infty)$
  - I will try to set as much as possible in discrete time because it's easier
- A is the menu; typically binary $A = \{\ell, r\}$
- For each s the analyst observes $\rho^s \in \Delta(A \times \mathcal{T})$

Example: character recognition

- $A = \{c, e\}$ is the menu
- $S = \{s^c, s^e\}$ is the true character
- Analyst knows the true s and runs the experiment many times for each s to collect empirical frequencies
- $\rho^s(c, t)$ probability that subject decides for exactly t seconds and chooses c if the true character is s

Example: weight discrimination

- $A = \{\ell, r\}$ is the menu
- $s = (s_\ell, s_r)$ is the true weight of each item, so $S = \mathbb{R}^2_+$
- Analyst knows the true s and runs the experiment many times for each s to collect empirical frequencies
- $\rho^s(\ell, t)$ probability that subject decides for exactly t seconds and chooses ℓ if the true weights are s

General Model

- At each time t the agent receives a message $m_t \in M_t$
- $m^t := (m_1, \ldots, m_t)$ denotes the history of messages up to time t
- The agent has a prior $p \in \Delta(S)$ and a utility $v : S \to \mathbb{R}^X$
- If forced at t, choice is $\chi_t = x$ iff $\mathbb{E}[v(x)|m^t] = \max_{y \in A} \mathbb{E}[v(y)|m^t]$
  - this is exactly our static BEU model from last lecture
- But the agent can always delay and get more signals (at a cost)

Stopping Time

Key idea: stopping at time t depends only on messages up to time t

Formally:

- Useful to think of the big probability space $\Omega = S \times \left(\bigtimes_{t \in \mathcal{T}} M_t \right)$
- $\mathbb{P} \in \Delta(\Omega)$ formed using the prior on S and the conditionals over M_t
- For any $\omega = (s, m_1, m_2, \ldots)$ we will denote $m^t(\omega) := (m_1, \ldots, m_t)$
- For each t there is a *stopping region* $\Sigma^t \subseteq M^t$

Definition A *stopping time* τ is a mapping $\tau : \Omega \to \mathcal{T}$ such that for each t we have $\tau(\omega) = t$ iff $m^t(\omega) \in \Sigma^t$.

Optimal Stopping

- Cost of waiting, a deterministic non-decreasing function $C : \mathcal{T} \to \mathbb{R}_+$

- The *optimal stopping* time τ^* solves:

$$\max_{\tau} \mathbb{E}[v(\chi_\tau) - C(\tau)]$$

- In statistics, this is known as *sequential sampling*: the analyst can buy additional data (experiments) at a cost.

- The special case of *linear* time cost is often used where $C(t) = ct$ for some $c > 0$.

Wald's Model

- Linear time cost; binary menu $A = \{\ell, r\}$
- Two states $S = \{s^\ell, s^r\}$
- Payoffs $v(x, s) = \mathbb{1}_{\{s = s^x\}}$
- Conditional on s, messages are i.i.d. $m_t \sim \mathcal{N}(\delta(s), \sigma^2)$, where $\delta(s^\ell) = d$ and $\delta(s^r) = -d$
- It is sufficient for the agent to keep track of the running sum $\bar{m}^t := m_1 + \cdots + m_t$, instead of the whole vector m^t
- $\bar{m}^t$ is a random walk with unknown drift (d or $-d$) that the agent is learning about. By Bayes rule, the posterior log-likelihood ratio is

$$\log \frac{\mathbb{P}(s^\ell | \bar{m}^t)}{\mathbb{P}(s^r | \bar{m}^t)} = \log \frac{\mathbb{P}(s^\ell)}{\mathbb{P}(s^r)} + \bar{m}^t \frac{2d}{\sigma^2}$$

Wald's Model—forced stopping

- The posterior log-likelihood ratio is

$$L_t := \log \frac{\mathbb{P}(s^\ell | \bar{m}^t)}{\mathbb{P}(s^r | \bar{m}^t)} = \log \frac{\mathbb{P}(s^\ell)}{\mathbb{P}(s^r)} + \bar{m}^t \frac{2d}{\sigma^2}$$

- If forced at time t the agent picks ℓ whenever $L_t > 0$
- For symmetric prior $L_t > 0$ iff $\bar{m}^t > 0$
- In state s_ℓ, at time t the agent chooses ℓ with probability $\mathbb{P}^{s^\ell}(\bar{m}^t > 0) = 1 - \Phi\left(\frac{-td}{\sigma\sqrt{t}}\right)$, where Φ is the cdf of $N(0,1)$
- This function is increasing in t, which formalizes the intuitive reasoning behind the speed-accuracy tradeoff

Wald's Model—optimal stopping

- In the Wald model this speed-accuracy tradeoff is *exactly* offset by optimal stopping
- On balance, accuracy is a constant function of time!
    - $\rightsquigarrow$ the reason for this will become clear in a couple of slides

Wald's Model—optimal stopping

Theorem: In the Wald model there exists $k > 0$ such that

$$\tau^* = \min\{t \geq 0 : |L_t| \geq k\},$$

Moreover, if the prior is symmetric, τ^* can also be written as

$$\tau^* = \min\{t \geq 0 : |\bar{m}^t| \geq b\}$$

for some $b > 0$.

The Wald model

Theorem: With symmetric prior the optimal strategy in the Wald model is

$$\tau^* := \min\{t \geq 0 : |\bar{m}^t| \geq b\} \qquad \chi_\tau := \begin{cases} \ell & \text{if} \quad \bar{m}^\tau = b \\ r & \text{if} \quad \bar{m}^\tau = -b \end{cases}$$

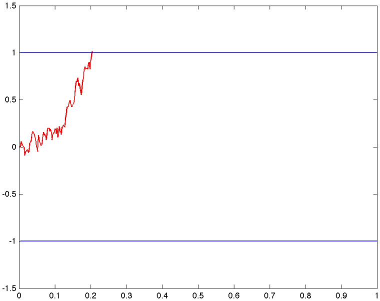

The Wald model

Theorem: With symmetric prior the optimal strategy in the Wald model is

$$\tau^* := \min\{t \geq 0 : |\bar{m}^t| \geq b\} \qquad \chi_\tau := \begin{cases} \ell & \text{if} \quad \bar{m}^\tau = b \\ r & \text{if} \quad \bar{m}^\tau = -b \end{cases}$$

The Wald model

Theorem: With symmetric prior the optimal strategy in the Wald model is

$$\tau^* := \min\{t \geq 0 : |\bar{m}^t| \geq b\} \qquad \chi_\tau := \begin{cases} \ell & \text{if} \quad \bar{m}^\tau = b \\ r & \text{if} \quad \bar{m}^\tau = -b \end{cases}$$

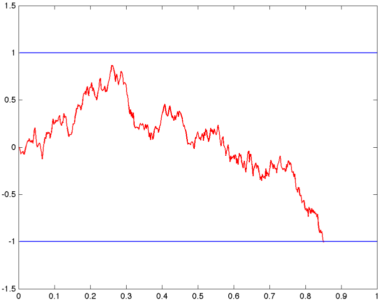

Comments

- Brought to the psychology literature in the 1960s and 1970s to study perception and memory retrieval
- Used extensively; well established in psych and neuroscience
- Ofen people *abstract from the optimization problem* and use this solution as a reduced-form model to generate $\rho \in \Delta(A \times \mathcal{T})$
- A continuous-time version of this reduced-form model is called the *Drift-Diffusion Model* (DDM)

DDM

Definition: Fix $A = \{\ell, r\}$. $\rho \in \Delta(A \times \mathcal{T})$ has a *DDM representation* if there exists $\delta \in \mathbb{R}$ and $\sigma, b > 0$ such that the cumulative signal is a *diffusion*

$$\bar{m}^t = t\delta + \sigma B_t,$$

where $\delta \in \mathbb{R}$ is the *drift* and B_t is a standard Brownian motion and ρ is the joint distribution induced by τ and χ, where

$$\tau = \inf \{t \geq 0 : |\bar{m}^t| \geq b\},$$
$$\chi_t = \ell \text{ iff } \bar{m}^t \geq b.$$

Notation: In this case we write $\rho \sim DDM(\delta, \sigma, b)$

Connection to Wald: $\rho^{s^\ell} \sim DDM(d, \sigma, b)$ and $\rho^{s^r} \sim DDM(-d, \sigma, b)$

Gambler's ruin problem

Theorem: If $\rho \sim DDM(\delta, \sigma, b)$, then

- the parameters are unique up to a common positive scalar multiple
- ρ is a product measure over $A \times \mathcal{T}$, i.e., accuracy is constant over time
- for any $t \in \mathcal{T}$ the conditional choice probability equals

$$\rho(\ell) = \frac{e^{\delta b/\sigma^2}}{e^{\delta b/\sigma^2} + e^{-\delta b/\sigma^2}}$$

and

$$\mathbb{E}[\tau] = \frac{b}{\delta} \tanh\left(\frac{b\delta}{\sigma^2}\right),$$

where tanh is the hyperbolic tangent function; $\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$.

Psychometric Function and Chronometric Function

- If we look at $\rho(\ell)$ as a function of $\delta \rightsquigarrow$ *psychometric function*
- If we look at $\mathbb{E}[\tau]$ as a function of $\delta \rightsquigarrow$ *chronometric function*

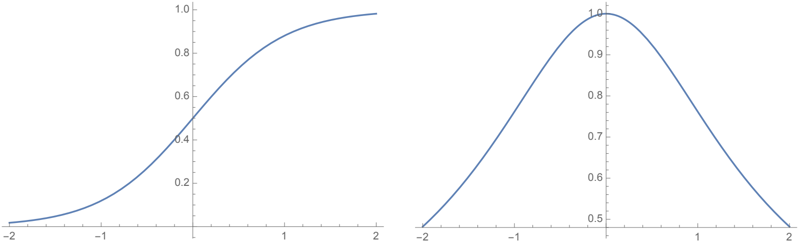

Figure: The psychometric (left) and chronometric functions (right). Here δ varies over the interval $[-2, 2]$ and $b = \sigma = 1$.

Basic Problem #1 with DDM

- In the Wald model there were two states $\delta = +d$ or $\delta = -d$ for some fixed d
  - well suited to tasks like character recognition (two characters)
  - Wald's theorem said DDM was the optimal thing to do there
- But now we seem to have a continuum of states $\delta \in [-2, 2]$
  - corresponds to an experiment where there many possible weights
  - indeed, DDM often applied to weight discrimination tasks and the like
  - no theorem says DDM is the optimal thing to do here!
  - this is a different learning problem: agent is learning about the *intensity* of the stimulus as well as the sign

Basic Problem #2 with DDM

- DDM predicts constant accuracy, while the stylized fact is that accuracy is decreasing

- Tweaks of DDM have been proposed to address that:

  – "full DDM"/"extended DDM": randomize over: 1) the drift, 2) the starting point of $\bar{m}^t$, and 3) the initial latency (non-response period)

    – this seems really ad-hoc!

  – "accumulator Models" or "race models": each item has its own signal accumulation process and its own boundary

    – contrast with DDM where the boundary is on the difference
    – is this ad hoc or microfounded?

  – time-dependent DDM: make the boundary a function of time

    – we will see this actually has a microfoundation

time-dependent DDM

Definition: Fix $A = \{\ell, r\}$. The s.c.f. $\rho \in \Delta(A \times \mathcal{T})$ has a *time-dependent DDM representation* if there exists $\delta \in \mathbb{R}$ and $\sigma > 0, b : \mathcal{T} \to \mathbb{R}_+$ such that the cumulative signal is a *diffusion*

$$\bar{m}^t = t\delta + \sigma B_t,$$

where $\delta \in \mathbb{R}$ is the drift and B_t is a standard Brownian motion and ρ is the joint distribution induced by τ and χ, where

$$\tau = \inf \{t \geq 0 : |\bar{m}^t| \geq b(t)\},$$
$$\chi_t = \ell \text{ iff } \bar{m}^t \geq b(t).$$

Notation: In this case we write $\rho \sim DDM^+(\delta, \sigma, b)$.

DDM^{+}

$$\tau = \inf \{t \geq 0 : |\bar{m}^t| \geq b(t)\}$$

$$\chi_t = \ell \text{ iff } \bar{m}^t \geq b(t)$$

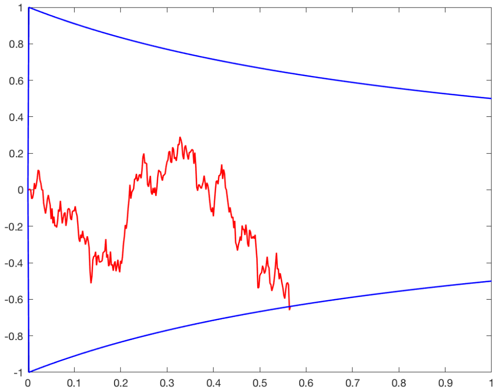

DDM^+

$$\tau = \inf \{t \geq 0 : |\bar{m}^t| \geq b(t)\}$$

$$\chi_t = \ell \text{ iff } \bar{m}^t \geq b(t)$$

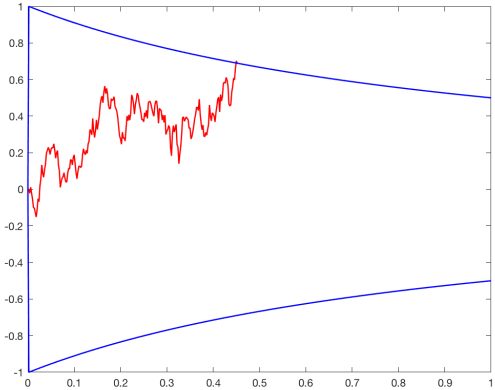

DDM^+

$$\tau = \inf \{t \geq 0 : |\bar{m}^t| \geq b(t)\}$$
$$\chi_t = \ell \text{ iff } \bar{m}^t \geq b(t)$$

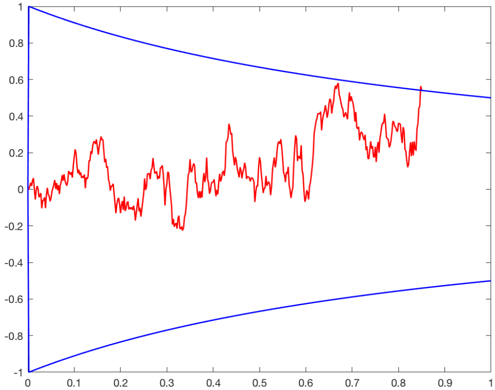

DDM^+

Theorem Suppose that $\rho \sim DDM^+(\delta, \sigma, b)$.

$$\text{Accuracy is } \begin{Bmatrix} \text{increasing} \\ \text{decreasing} \\ \text{constant} \end{Bmatrix} \text{ iff boundary } b \text{ is } \begin{Bmatrix} \text{increasing} \\ \text{decreasing} \\ \text{constant} \end{Bmatrix}$$

Intuition for decreasing accuracy: higher bar to clear for small t, so if the agent stopped early, $\bar{m}^t$ must have been high, so higher likelihood of making the correct choice

Microfounding a time-dependent boundary

- So far, only the constant boundary b was microfounded
- Do any other boundaries come from optimization?
- What is the optimization problem?
- We now derive the optimal boundary

Chernoff's Model

- Linear time cost; binary menu $A = \{\ell, r\}$
- Continuum of states $S = \mathbb{R}^2$; $s = (s_\ell, s_r)$. We have $v(x, s) = s_x$
- Conditional on s, $m_{t,x} \sim^{i.i.d.} \mathcal{N}(s_x, \sigma^2)$ independent over $x \in A$
- The prior is $s_x \sim N(\mu_{0,x}, \sigma_0^2)$ independent over $x \in A$
- Sufficient to keep track of the running sum $\bar{m}_x^t := m_{1,x} + \cdots + m_{t,x}$
- $\bar{m}_x^t$ is a random walk with unknown drift. By Bayes rule, the posterior is $s_x \sim N(\mu_{t,x}, \sigma_t^2)$, where

$$\mu_{t,x} = \mu_{0,x} \frac{\sigma_t^2}{\sigma_0^2} + \bar{m}_x^t \frac{\sigma_t^2}{\sigma^2} \quad \text{and} \quad \sigma_t^{-2} = \sigma_0^{-2} + t\sigma^{-2}$$

Chernoff's Model

Theorem: In the Chernoff model there exists a decreasing function $k : \mathcal{T} \to \mathbb{R}$ such that

$$\tau^* = \inf\{t \geq 0 : |\mu_t| \geq k(t)\},$$

where $\mu_t := \mu_{t,\ell} - \mu_{t,r}$ is the posterior mean difference.

Moreover, if $\mu_0 = 0$, then there exists $b : \mathcal{T} \to \mathbb{R}$ such that

$$\tau^* = \inf\{t \geq 0 : |\bar{m}^t| \geq b(t)\},$$

where $\bar{m}^t = \bar{m}^t_\ell - \bar{m}^t_r$.

Corollary: In Chernoff's model $\rho^s \sim DDM^+(s_\ell - s_r, \sigma\sqrt{2}, b)$.

Key difference between Wald and Chernoff

- Intuition for Wald: **stationarity**
  - suppose that you observe $\bar{m}_\ell^t \approx \bar{m}_r^t$ after a long t
  - you know drift cannot be zero
  - you think to yourself: "the signal must have been noisy"
  - so you don't learn anything $\Rightarrow$ you continue

- Intuition for Chernoff: **non-stationarity**
  - suppose that you observe $\bar{m}_\ell^t \approx \bar{m}_r^t$ after a long t
  - you think to yourself: "I must be indifferent"
  - so you have learned a lot $\Rightarrow$ you stop

- Intuition for the difference between the two models:
  - interpretation of signal depends on the prior

A different model for perception

- Chernoff model is good for economic decisions: $v(x, s) = s_x$
  - you get the utility of what you consume
- A model for perception would have $v(x, s) = \mathbb{1}_{s_x > s_y}$
  - reward independent of how hard the choice is
  - this model also leads to DDM^+ but with a different boundary

Is any boundary optimal?

Theorem: For any b there exists a (nonlinear) cost function C such that b is the optimal solution in the Chernoff model

Do difficult choices take more time?

- Mechanically true in DDM,
  - harder choice = $|\delta|$ smaller
  - chronometric function is hump-shaped around zero

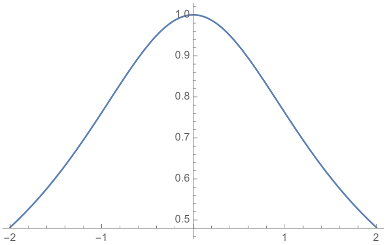

Do difficult choices take more time?

- Actually, this is true in all DDM$^+$,
  - harder to show
- But what is the intuition? why spend more time if almost indifferent?
  - if knew that indifferent, just toss a coin and spend zero time
  - but you don't know you are almost indifferent—start with your prior!
  - once you learn you are indifferent, then stop