Speed, Accuracy, and the Optimal Timing of Choices

Drew Fudenberg, MIT

Philipp Strack, Berkeley

Tomasz Strzalecki, Harvard

Choices

Suppose we see an agent's choices from the menu $A = \{\ell, r\}$.

Sometimes she chooses ℓ, sometimes she chooses r.

How to model?

- **choice correspondence**: the agent is indifferent, regardless of the relative probabilities
- **stochastic choice function**: treat choice probabilities as data

Decision Times

New Variable: How long does the agent take to decide?

Time: $\mathcal{T} = [0, \infty)$

Observe: Joint distribution $P \in \Delta(A \times \mathcal{T})$

Question:

- Are fast decisions “better” or “worse” than slow ones?

Behavioral Story

System I is instinctive and fast, System II is deliberative and slow, so fast decisions are worse (Kahneman)

- we will ignore this and have a one-system story

Are quick decisions better than slow ones?

Informational Effect:

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

Selection Effect:

- Time is costly, so you decide to stop depending on how much you expect to learn (option value of waiting)
  - Want to stop early if get an informative signal
  - Want to continue if get a noisy signal

- This creates dynamic selection
  - stop early after informative signals
  - informative signals more likely when the problem is easy

Decreasing accuracy

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

Stylized fact: Decreasing accuracy: fast decisions are "better"

- Well established in perceptual tasks (dots moving on the screen), where "better" is objective
- Also in experiments where subjects choose between consumption items

When are decisions "more accurate?"

In cognitive tasks, **accurate = correct**

In choice tasks, **accurate = preferred**

$\boldsymbol{p(t)}$:= probability of making the correct/preferred choicce conditional conditional on deciding at t

Definition:

$$P \text{ displays } \left\{\begin{matrix}\text{increasing}\\ \text{decreasing}\\ \text{constant}\end{matrix}\right\} \text{ accuracy iff } p(t) \text{ is } \left\{\begin{matrix}\text{increasing}\\ \text{decreasing}\\ \text{constant}\end{matrix}\right\}$$

Experiment of Krajbich, Armel, and Rangel (2010)

- X: 70 different food items
- Step 1: Rate each $x \in X$ on the scale -10, . . . , 10
- Step 2: Choose from $A = \{\ell, r\}$ (100 different pairs)
  - record choice and decision time
- Step 3: Draw a random pair and get your choice

Decreasing Accuracy

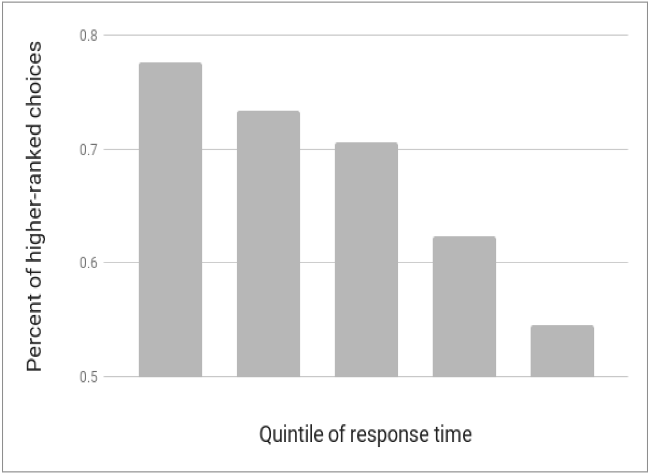

based on data from Krajbich, Armel, and Rangel (2010)

Related Literature: Economics

* *stochastic choice conditional on exogenous time*
Caplin and Dean (2011); Natenzon (2014); Lu (2016);
Cerreia-Vioglio, Maccheroni, Marinacci, and Rustichini (2017)

* *stochastic choice and endogenous timing of decisions*
Che and Mierendorff (2016); Hebert and Woodford (2017);
Jehiel and Steiner (2017); Woodford (2014);

Related Literature

* Cognitive science: *drift diffusion model (DDM)*

  e.g. Ratcliff-McKoon (2008); Krajbich et al (2010), Milosavljevic et al (2010); Drugowitsch et al (2012); Clithero- Rangel (2014)

* Probability and statistics: *optimal stopping*

  e.g. Wald (1947); Arrow, Blackwell, Girshick (1949); Chernoff (1961); Bather (1962); Peskir and Shiryaev (2006)

* Less closely related: *classify decisions as "instinctive/heuristic" or "cognitive"*

  e.g. Rubinstein (2007); Kahneman (2011); Rand et al (2012); Caplin and Martin (2015)

Model

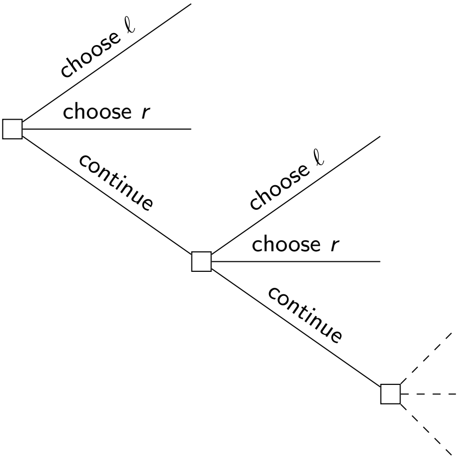

Learning Model

- unknown utility $\theta = (\theta^\ell, \theta^r) \in \mathbb{R}^2$; prior belief on θ

- observe a two-dimensional signal for $i = \ell, r$

  $Z_t^i = \theta^i t + \alpha B_t^i$
  B_t^i are independent Brownian motions
  let $Z_t := Z_t^\ell - Z_t^r$

- examples of prior/posterior families

  - "certain difference"

    – binomial prior: either $\theta = (1, 0)$ or $\theta = (0, 1)$
    – binomial posterior: either $\theta = (1, 0)$ or $\theta = (0, 1)$

  - "uncertain difference"

    – Gaussian prior: $\theta^i \sim N(X_0^i, \sigma_0^2)$
    – Gaussian posterior: $\theta^i \sim N(X_t^i, \sigma_t^2)$

Interpretation of the Signal Process

- recognition of the objects on the screen
- retrieving pleasant or unpleasant memories
- coming up with reasons pro and con
- introspection
- signal strength depends on the utility difference or on the ease of the perceptual task

In animal experiments, some neuroscientists record neural firing and relate it to these signals

We don't do this, treat signals as unobserved by the analyst

Learning Model

- τ is a stopping time (measurable w.r.t Z_t)
- conditional on stopping, the agent maximizes expected utility

$$\text{choice}_\tau = \text{argmax}\{\mathbb{E}_\tau \theta^\ell, \mathbb{E}_\tau \theta^r\}$$

Example: If stopping is exogenous (τ is independent of signal Z_t), and prior is symmetric, there is **increasing** accuracy: waiting longer gives better information so generates better decisions

Exogenous vs Endogenous Stopping

- Key assumption above: stopping independent of signal
- If stopping is conditional on the signal, this could get reversed
- Intuition: with endogenous stopping you
  - *#1* stop early after informative signals (and make the right choice); wait longer after noisy signals (and possibly make a mistake)
  - *#2* probably faced an easier problem if you decided quickly

Optimal Stopping

The agent chooses a Z_t-measurable stopping time τ to optimize:

$$\max_{\tau} \Big[\max\{\mathbb{E}_\tau \theta^l, \mathbb{E}_\tau \theta^r\} - c\tau \Big]$$

(we focus on the "minimal optimal" stopping time)

The "certain difference" model

* Assumptions:

- binomial prior: either $\theta = (1, 0)$ or $\theta = (0, 1)$
- binomial posterior: either $\theta = (1, 0)$ or $\theta = (0, 1)$

* Key intuition: **stationarity**

- suppose that you observe $Z_t^l \approx Z_t^r$ after a long t
- you think to yourself: "the signal must have been noisy"
- so you don't learn anything $\Rightarrow$ you continue

* Formally, the option value is constant in time

The "certain difference" model

Theorem: (Wald, Arrow, Blackwell, Girshick, Shiryaev)
When the prior is symmetric, the optimal stopping time is

$$\tau^* = \inf\{t \geq 0 \colon |Z_t| \geq b\}$$

where $b > 0$.

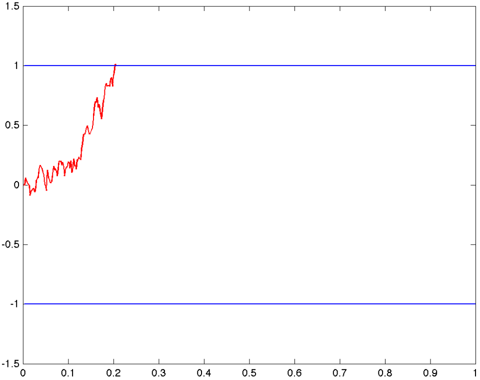

$$\tau^* = \inf\{t \geq 0 \colon |Z_t| \geq b\}$$

$$\tau^* = \inf\{t \geq 0\colon |Z_t| \geq b\}$$

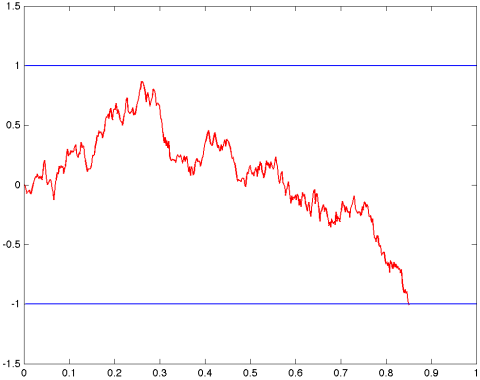

$$\tau^* = \inf\{t \geq 0\colon |Z_t| \geq b\}$$

Hitting Time Models

- can use this algorithm to generate a distribution $P \in \Delta(A \times \mathcal{T})$ without worrying about optimality

- closed forms for choice probabilities and mean stopping time

- used extensively for perception tasks since the 70's; pretty well established in psych and neuroscience

- more recently used to study choice tasks by a number of teams of authors including Colin Camerer and Antonio Rangel

- Many versions of the model
  - ad-hoc tweaks (not worrying about optimality)
    - assumptions about the process Z_t
    - functional forms for the time-dependent boundary
  - much less often, optimization used:
    - time-varying costs (Drugovitsch et al, 2012)
    - endogenous attention (Woodford, 2014)

Hitting Time Models

* **Definition**:

  - stochastic process Z_t starts at 0
  - time-dependent boundary $b : \mathbb{R}_+ \to \mathbb{R}_+$
  - hitting time $\tau = \inf\{t \geq 0 : |Z_t| \geq b(t)\}$
  - choice $= \begin{cases} \ell & \text{if} \quad Z_\tau = +b(\tau) \\ r & \text{if} \quad Z_\tau = -b(\tau) \end{cases}$

Drift Diffusion Model (DDM)

Special case where the process Z_t is a diffusion with constant drift and volatility

$$Z_t = \delta t + \alpha B_t$$

(could eliminate the parameter α here but it's useful later)

Definition: P has a **DDM representation** if it can be represented by a stimulus process $Z_t = \delta t + \alpha B_t$ and a time-dependent boundary b. We write this as $P = P(\delta, \alpha, b)$.

average DDM

Definition: P has an **average DDM representation** $P(\mu, \alpha, b)$ with $\mu \in \Delta(\mathbb{R})$ if $P = \int P(\delta, \alpha, b) d\mu(\delta)$.

- in an average DDM model the analyst does not know δ, but has a correct prior
- intuitively, it is unknown to the analyst how hard the problem is for the agent

Hitting Time Model

Proposition: Any Borel $P \in \Delta(A \times T)$ has a hitting time representation where the stochastic process Z_t is a time-inhomogeneous Markov process and the barrier is constant

Remarks:

- this means that the general model is without loss of generality
- in particular, it is without loss of content to assume that b is independent of time
- however, in the general model the process Z_t may have jumps
- From now on we focus on the DDM special cases

DDM

Definition: Accuracy in DDM is the probability of making the choice which agrees with the signal

$$p(t) = \mathbb{P}\left[\operatorname{sgn} Z_\tau = \operatorname{sgn} \delta \mid \tau = t\right].$$

- In DDM p is the probability of making the modal choice.
- If the correct choice is part of the data, this is the probability of making the correct choice

Accuracy in DDM models

Theorem: Suppose that $P = P(\delta, \alpha, b)$.

$$P \text{ displays } \begin{Bmatrix} \text{increasing} \\ \text{decreasing} \\ \text{constant} \end{Bmatrix} \text{ accuracy iff } b \text{ is } \begin{Bmatrix} \text{increasing} \\ \text{decreasing} \\ \text{constant} \end{Bmatrix}$$

Intuition for decreasing accuracy: this is our selection effect #1

- higher bar to clear for small t, so if the agent stopped early, Z must have been very high, so higher likelihood of making the correct choice

Accuracy in DDM models

Theorem: Suppose that $P = P(\mu, \alpha, b)$, with $\mu = \mathcal{N}(0, \sigma_0)$

$$P \text{ displays } \begin{Bmatrix} \text{increasing} \\ \text{decreasing} \\ \text{constant} \end{Bmatrix} \text{ accuracy iff } b(t) \cdot \sigma_t \text{ is } \begin{Bmatrix} \text{increasing} \\ \text{decreasing} \\ \text{constant} \end{Bmatrix}$$

where $\sigma_t^2 := \frac{1}{\sigma_0^{-2} + \alpha^{-2} t}$

Intuition for decreasing accuracy: this is our selection effect #2

- σ_t is a decreasing function; this makes it an easier bar to pass

selection effect #2

Proposition: Suppose that $\mu = \mathcal{N}(0, \sigma_0)$, and $b(t) \cdot \sigma_t$ non-increasing. Then $|\delta|$ **decreases in** τ in the sense of FOSD, i.e. for all $d > 0$ and $0 < t < t'$

$$\mathbb{P}\left[\,|\delta| \geq d \mid \tau = t\right] > \mathbb{P}\left[\,|\delta| \geq d \mid \tau = t'\right] .$$

- larger values of $|\delta|$ more likely when the agent decides quicker
- problem more likely to be "easy" when a quick decision is observed
- this is a selection coming from the analyst not knowing how hard the problem is

Microfounding the Boundary

* So far, only the constant boundary b was microfounded

* Do any other boundaries come from optimization?

* Which boundaries should we use?

* We now derive the optimal boundary

The "uncertain difference" model

* Assumptions:
  - Gaussian prior: $\theta^i \sim N(X_0^i, \sigma_0^2)$
  - Gaussian posterior: $\theta^i \sim N(X_t^i, \sigma_t^2)$

* Key intuition: **nonstationarity**
  - suppose that you observe $Z_t^l \approx Z_t^r$ after a long t
  - you think to yourself: "I must be indifferent"
  - so you have learned a lot $\Rightarrow$ you stop

* Formally $\sigma_t^2 = \frac{1}{\sigma_0^{-2} + \alpha^{-2} t}$ so option value is decreasing in time

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

The "uncertain difference" model

Theorem:

1. There is a strictly decreasing, strictly positive $k^* : \mathbb{R}_+ \to \mathbb{R}_+$ such that

$$\tau^* = \inf\{t \geq 0\colon |X_t^l - X_t^r| \geq k^*(t)\}.$$

Moreover $\lim_{t\to\infty} k_t^* = 0$.

2. If $X_0^l = X_0^r$, there is a strictly positive $b^* : \mathbb{R}_+ \to \mathbb{R}_+$ such that

$$\tau^* = \inf\{t \geq 0\colon |Z_t^l - Z_t^r| \geq b^*(t)\},$$

where $b^*(t) = \alpha^2 \sigma_t^{-2} k^*(t)$. Furthermore, we have the following bounds on the slope of b^*

$$-b^*(t)\sigma_t^2 \;\leq\; b^{*\prime}(t) \;\leq\; \frac{1}{2} b^*(t)\sigma_t^2$$

Part 1. follows from the principle of optimality for continuous time processes and the shift invariance property of the value function, which is due to the normality of the posterior.

Part 2. describes the optimal strategy τ^* in terms of stopping regions for the signal process $Z_t := Z_t^l - Z_t^r$. This facilitates comparisons with the simple DDM, where the process of beliefs lives in a different space and is not directly comparable.

Intuitions

- k^* strictly decreasing because belief updating slows down
- k^* decreases all the way to 0 because otherwise the agent would have a positive subjective probability of never stopping and incurring an infinite cost
  - note: in the simple DDM, the agent is sure that the absolute value of the drift of the signal is bounded away from 0, so she believes she will stop in finite time with probability 1 even though the boundaries are constant

Proposition: The average "uncertain difference" DDM has decreasing accuracy, i.e. the probability that the agent makes the correct choice

$$\mathbb{P}\left[\operatorname{sgn}\left(X^{l}_{\tau^*} - X^{r}_{\tau^*}\right) = \operatorname{sgn}\delta \mid \tau^* = t\right]$$

decreases in t.

Endogenous Attention

- The agent can choose attention levels

$$\beta_t^l, \beta_t^r \geq 0$$

- Attention influences the signals Z_t^1, Z_t^2

$$dZ_t^i = \beta_t^i \theta^i dt + dB_t^i .$$

- Fixed attention budget $\beta_t^l + \beta_t^r \leq 2$
- $\beta_t^l = \beta_t^r = 1$ leads to the same signal process as before
- $\alpha = 1$ for simplicity here

Endogenous Attention

Theorem: The optimal attention strategy pays equal attention to both signals

$$\beta_t^l = \beta_t^r = 1$$

and thus leads to the same choice process as the exogenous attention model.

Endogenous Attention

Theorem: The optimal attention strategy pays equal attention to both signals

$$\beta_t^l = \beta_t^r = 1$$

and thus leads to the same choice process as the exogenous attention model.

Intuition: This strategy minimizes the posterior variance of the difference in posterior means $X_t^l - X_t^r$ at *every* point in time t and thus maximizes the speed of learning.

The Chernoff (1961) model

Regret Minimization: for any stopping time τ the objective function is

$$\mathrm{Ch}\,(\tau) := \mathbb{E}\left[-\mathbf{1}_{\{x_\tau^l \geq x_\tau^r\}}(\theta^r - \theta^l)^+ - \mathbf{1}_{\{x_\tau^r > x_\tau^l\}}(\theta^l - \theta^r)^+ - c\tau\right]$$

the agent gets zero for making the correct choice and is penalized the foregone utility for making the wrong choice

The Chernoff (1961) model

Theorem: For any stopping time τ

$$\mathrm{Ch}\,(\tau) = \mathbb{E}\left[\max\{X^l_\tau, X^r_\tau\} - c\tau\right] + \kappa,$$

where κ is a constant independent of τ; therefore, these two objective functions induce the same choice process.

The Chernoff (1961) model

Theorem: For any stopping time τ

$$\mathrm{Ch}\,(\tau) = \mathbb{E}\left[\max\{X^l_\tau, X^r_\tau\} - c\tau\right] + \kappa,$$

where κ is a constant independent of τ; therefore, these two objective functions induce the same choice process.

Intuition: Subtracting the expected value of the optimal choice $\mathbb{E}[\max\{\theta^l, \theta^r\}]$, using that τ is a stopping time and applying the law of iterated expectations conditional on either choice being correct yields the result.

Non-linear cost

Theorem: Consider either the Certain or the Uncertain-Difference DDM. For any finite boundary b and any finite set $G \subseteq \mathbb{R}_+$ there exists a cost function $d : \mathbb{R}_+ \to \mathbb{R}$ such that b is optimal in the set of stopping times T that stop in G with probability one

$$\inf\{t \in G \colon |Z_t| \geq b(t)\} \in \text{argmax}_{\tau \in \mathcal{T}} \mathbb{E}\left[\max\{X^1_\tau, X^2_\tau\} - d(\tau)\right] .$$

Application/Experiment

- data of Krajbich, Armel and Rangel (2010)
- 39 subjects making choices between food items
- asked to refrain from eating for 3 hours before the experiment
- each subject asked to make 100 pairwise choices
- also separately elicited ratings of these items (on the scale from -10 to $+10$)

Fitting to a closed-form boundary

- We consider two functional forms:
  - $b(t) = \frac{1}{g+ht}$ which is the approximately optimal boundary
  - $b(t) = g\exp(-ht)$ used in Milosavljevic et al. (2010)
- In any case, the parameters are (δ, α, g, h)
  - δ is the drift; we take it to be the difference in the numerical ratings of the two items in the choice set
  - α is the volatility of Z_t
  - (g, h) are the parameters of the boundary

Fitting to a closed-form boundary

- For each (δ, α, g, h) need to compute the joint probability density of stopping and choice (the likelihood function)
  - to compute the distribution of hitting times, we used Monte Carlo simulations with 1 million random paths (this takes about a week on a cluster)
  - the conditional choice probabilities as a function of stopping time are given in closed form
- Then use the gradient descent algorithm to find maximum
- **Findings**: for 30 out of 39 subjects the approximately optimal boundary is a better fit than the exponential boundary

Fitting to the optimal boundary

- Additional computation: the optimal boundary

- We computed this by imposing a large finite terminal time, discretizing time and space on a fine grid, and solving backwards. This computation only needs to be done for a single parameter constellation, due to a result in the online appendix, and takes less than two hours on a laptop

- Then compute the likelihood function as above

- **Findings**:
    - there is substantial heterogeneity between the subjects
    - two out of 39 subjects have a non-monotone boundary

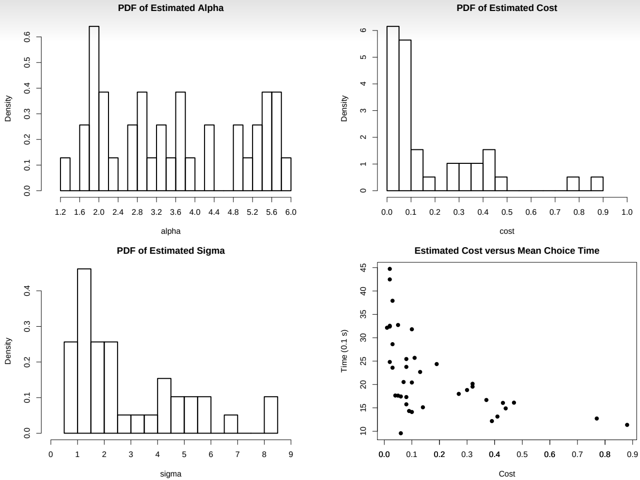

Figure: Marginal distributions of α, c, and σ_0 and the correlation between the average stopping time for each subject and their estimated cost c.

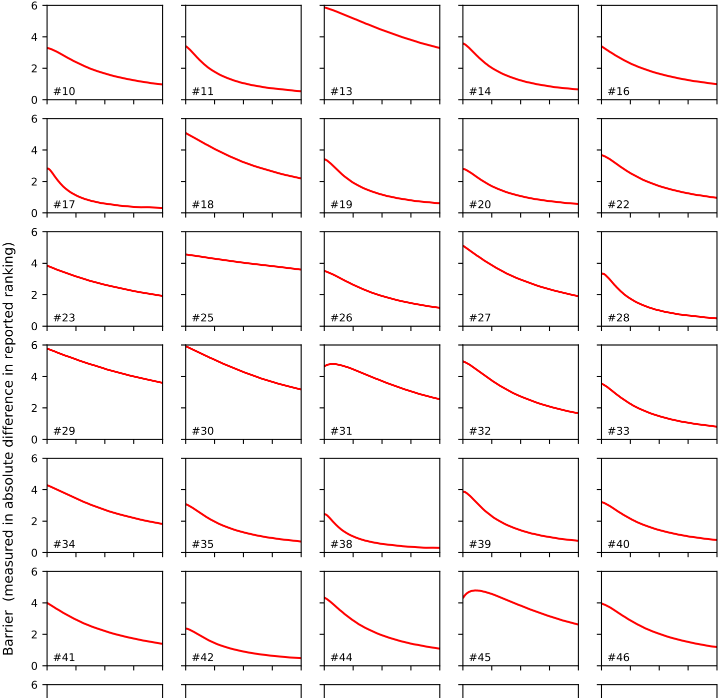

Barrier (measured in absolute difference in reported ranking)
0
2
4
6
#10
#11
#13
#14
#16
#17
#18
#19
#20
#22
#23
#25
#26
#27
#28
#29
#30
#31
#32
#33
#34
#35
#38
#39
#40
#41
#42
#44
#45
#46

Recap

- Observables: joint distribution over choices and decision times
- General DDM: Brownian signals and arbitrary boundary
  - characterize when earlier decisions better
- DDM derived from optimal stopping: Gaussian prior
  - allows agent to learn the choice is a toss-up
  - resulting boundary better fits the data than the constant boundary of simple DDM
  - explains why quicker choices are often better

Thank you!