

On Risk and Time Pressure: When to Think and When to Do

Christoph Carnehl Johannes Schneider

April 2022

Problem Solving under Time Pressure

An agent has to solve a problem by a deadline.

The agent can take two approaches:

Do: Implement an existing method.

Think: Think about an alternative method and implement it once found.

What is the difference between the two approaches?

- Doing is quick but carries fundamental risk. Doing is quick but carries fundamental risk.
- Thinking is fundamentally safer but carries time risk.
- Thinking is fundamentally safer but carries time risk.

How should a time-constrained agent allocate time between thinking and doing?

Entrepreneurial Problem Solving

Entrepreneurs have to reach milestones to secure follow-up funding.

- Implicit deadlines: running out of cash.
- Explicit deadlines: contract includes deadline.

An entrepreneur does if she attempts implementing her initial idea.

- launching a product in a market, improving a prototype, ...

An entrepreneur thinks if she considers a pivot—a substantial change in her strategy.

- customer research to find a target market, build a different product, ...

Pivots triggered by both approaching deadlines and low beliefs.

e.g.: Kirtley and O'Mahony, 2020; Rahmani and Ramachandran, 2021

Entrepreneurs face a time-risk tradeoff while working towards a milestone.

Not a Standard Bandit Problem

Doing: Exponential bandit that has unknown arrival rate.

Thinking: Exponential bandit with known arrival rate butbut triggers continuation game.

Economic setting is not a straightforward bandit problem:

- Thinking alone does not solve the problem.
- Methods that can solve the problem can endogenously expand.
- Finite horizon leads to time-dependent value of thinking.

Experimentation setting with endogenous arms and a deadline.

- Reduces to finite horizon two-armed bandit problem with one restless arm.

Time-risk tradeoff not captured in the classical infinite horizon experimentation framework (Rothschild, 1974; Weitzman, 1979).

Preview of Results

The agent's optimal policy will have at most three phases:

1. An initial doing phase.
2. A thinking phase.
3. A final doing, Hail Mary phase.

The initial doing phase reduces the probability of solving the problem but reduces the expected time worked.

- Rationalizes the false starts phenomenon in startups (Eisenmann, 2021).

Short deadlines can mitigate false starts and increase probability of successful thinking.

- Experimentation:
Rothschild (1974), Weitzman (1979), Bergemann and Välimäki (2008), Klein (2016), Fershtman and Pavan (2021), ...
- Strategy and Innovation:
Gans, Stern and Wu (2019), Felin, Gambardella, Stern and Zenger (2019), ...
- Thinking vs. Doing:
Bolton and Faure-Grimaud (2009), Kim (2021)
- Staged Experimentation:
Green and Taylor (2016), Moroni (2021), Wolf (2019)

Model

Model

Time is continuous, $t \in [0, T]$, with a deadline $T < \infty$.

The agent has access to a doing arm with unknown arrival rate.

- With $\bar{p} \in (0, 1)$ arrival rate $\lambda > 0$, otherwise 0.
- Absent an arrival belief declines according to Bayes' rule.
- Arrival generates payoff $B > 0$.

The agent has access to a thinking arm.

- With known arrival rate $\mu > 0$.
- Arrival at t generates continuation payoff $V(T - t)$.

Agent allocates effort at flow cost $c > 0$ on the two arms.

- $a_t \in [0, 1]$ on the initial idea (*doing*) with expected arrival rate $a_t p_t \lambda$.
- $1 - a_t \in [0, 1]$ on generating a new idea (*thinking*) with arrival rate $(1 - a_t) \mu$.

Assumptions on $V(T - t)$

New idea triggers continuation payoff that depends on time remaining $\tau := T - t$
→ new idea still needs implementation.

Assumptions on $V(\tau)$:

$V(0) = 0$ without time remaining not implementable

$V'(\tau) > 0$ more time increases implementation probability

$-\frac{V''(\tau)}{V'(\tau)} > \bar{p}\lambda$ sufficiently increasing relative time pressure

Note: Fundamentally different from a second arm with lower arrival rate. Payoff generated is time-dependent which alters incentives.

Benchmarks

Proposition

Suppose $T = \infty$. The agent either starts doing and switches to thinking eventually or the agent thinks throughout.

The agent pulls the doing arm if and only if

$$\bar{p} \geq \hat{p} := \frac{c/\lambda}{B - V(\infty) + c/\mu}.$$

The agent pulls the doing arm if the payoff of a good doing arm ($B - c/\lambda$) is higher than of an infinite horizon thinking arm ($V(\infty) - c/\mu$). Two potential sources:

- Higher value of a doing solution: $B > V(\infty)$
- Lower expected cost until arrival: $\lambda > \mu$

No Payoff Differences

Proposition

Suppose $c = 0$, $\lim_{\tau \rightarrow \infty} V(\tau) = B$ and $T < \infty$. The agent either starts thinking and switches to doing eventually or the agent never thinks.

The agent pulls the thinking arm if and only if the deadline is sufficiently long.

If the deadline is close, the value of thinking is low.

→ Agent will pull doing arm close to deadline.

Doing for a long time implies that agent will pull the doing arm while holding a low belief.

→ Agent will pull thinking arm with positive probability.

The thinking arm is most valuable when the deadline is far.

Why and When to Do?

Benchmarks provide four insights that guide the intuition for the optimal policy.

1. Close to the deadline, the agent has too little time to convert any progress on the thinking arm into a solution.
→ time pressure induces the agent to do late
2. The doing arm can require less effort when the agent is initially optimistic.
→ effort cost induce the agent to do early on
3. The doing arm becomes increasingly less valuable the more the agent thinks.
→ risk on the doing arm induces the agent to think
4. The thinking arm is most valuable when plenty of time remains.
→ time-dependence of thinking arm induces the agent to think early

Optimal Policy

The Agent's Problem

We relabel time from calendar time to time remaining, $\tau := T - t$.

Define as state variable

$$A_\tau := \int_\tau^T a_t dt : \text{effort allocated on doing before } \tau.$$

A_τ determines belief p_τ about doing arm.

Agent dynamically allocates effort between two arms to maximize payoff.

$$\max_{(a_\tau)_{\tau=0}^T} \int_0^T \underbrace{e^{-\mu(T-\tau-A_\tau)}}_{P(\text{no progress yet})} \underbrace{(1 - \bar{p} + \bar{p}e^{-\lambda A_\tau})}_{P(\text{no solution yet})} \underbrace{(\mu(1 - a_\tau)V(\tau) + \lambda a_\tau p_\tau B)}_{\text{flow payoff}} dt$$

Solve for optimal policy using optimal control and Pontryagin's maximum principle.

Characterization of the Optimal Policy

Proposition

An optimal policy exists. Absent the arrival of a success or a new idea, the optimal policy takes one of three forms

- 1. The agent exclusively uses the doing arm.*
- 2. The agent starts using the thinking arm and switches to the doing arm.*
- 3. The agent starts using the doing arm, switches to the thinking arm, and switches back to the doing arm.*

Intuitively:

- At the deadline, thinking has little value → doing at the end.

The Optimal Control Problem

The necessary conditions of the optimal control problem deliver a switching function that characterizes any candidate solution.

$$\gamma_\tau = e^{-\mu(T-\tau-A_\tau)} \left(1 - \bar{p} + \bar{p}e^{-\lambda A_\tau}\right) (\mu V(\tau) - p_\tau \lambda B) - \eta_\tau,$$

where η_τ is the co-state variable of the optimal control problem. [► Details](#)

If $\gamma_\tau < 0$, then the agent pulls the thinking arm with time τ remaining.

If $\gamma_\tau > 0$, then the agent pulls the doing arm with time τ remaining.

At the deadline, thinking has no value \rightarrow the agent pulls the doing arm.

$$\gamma_0 = -e^{-\mu(T-A_0)} \bar{p} e^{-\lambda A_0} \lambda B < 0$$

Evolution of the Switching Function

Structure of the optimal policy intuitively follows from evolution of switching function.

$$\dot{\gamma}_\tau \propto \left(\underbrace{\mu V'(\tau)}_{\text{(i) deadline effect}} + \underbrace{\mu p_\tau \lambda (V(\tau) - B)}_{\text{(ii) payoff-on-arrival effect}} + \underbrace{(\mu - \lambda p_\tau) c}_{\text{(iii) effort-to-arrival effect}} \right).$$

- (i) An increase in τ pushes the agent towards thinking.
- (ii) An increase in τ pushes the agent towards doing whenever $B > V(\tau)$.
- (iii) An increase in τ pushes the agent towards doing whenever $p_\tau \lambda > \mu$.

At Most Two Switches

Lemma

If the agent leaves the thinking arm, then she never returns to it.

The deadline effect pushes the agent always towards thinking ($V'(\tau) > 0$) but it becomes weaker ($V''(\tau) < 0$).

If an increase in the time remaining induced a switch *doing* \rightarrow *thinking*, then the deadline effect must have been relatively weak at that switch.

If the deadline effect weakens sufficiently relative to the increase in the thinking arm's value, $-V''(\tau)/V'(\tau) \geq \bar{p}\lambda$, the agent thinks at most once.

Note: Without relative concavity, potentially more switches due to payoff-on-arrival effect.

Characterization

The optimal policy can be described by

- τ_1 : length of initial doing period
- τ_2 : length of thinking period
- τ_3 : length of Hail Mary period

subject to

$$\tau_1 + \tau_2 + \tau_3 = T.$$


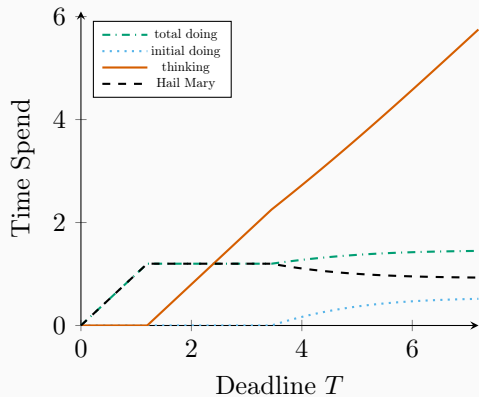
Proposition

Suppose $V(\infty) \leq B + c/\mu$ and well-behaved curvature of arms' payoffs if $\mu > \lambda$. There is a unique optimal policy that is found by a simple algorithm.

► Technical Condition

► Algorithm

Optimal Policy by Deadline



Entrepreneurial Problem Solving

Entrepreneurial Problem Solving

An entrepreneur has raised some funding for her venture but needs to reach a milestone to obtain a new round of financing in the future.

- Explicit deadlines and milestones frequent (e.g. Gompers, 1995)
- Implicit deadlines and milestones due to burn rate (e.g. failure due to running out of cash - CB Insight, 2021)

Entrepreneurs benefit from reaching the milestone early.

- More money in the pocket for later stages
- Action bias (Eisenmann, 2021)

Experimentation as classical innovation model but pivots—significant changes in strategy—have to be prepared before they can be implemented (Isenberg and DiFiore, 2020).

Examples

Examples of Thinking Value, I

An arrival on the thinking arm reveals a new method to launch the product.

A new exponential bandit arm with known arrival rate $\nu > \bar{p}\lambda$.

A success on this arm has value B_ν .

Pulling this arm has a flow cost of c_ν .

Examples of Thinking Value, II

An arrival on the thinking arm reveals a new method to launch the product.

A new bandit arm with unknown arrival rate.

With \bar{p}^ν , the arrival rate is $\nu > 0$ and 0 otherwise.

A success on this arm has value B_ν .

Pulling this arm has a flow cost of c_ν .

Examples of Thinking Value, III

An arrival on the thinking arm reveals a new method to launch the product.

A new bandit arm with time-varying intensity rate $\nu(\tau)$.

A success on this arm has value B_ν .

Pulling this arm has a flow cost of c_ν .

Time-varying intensity rate captures

- Learning-by-doing on a new approach (increasing $\nu(\tau)$).
- Solutions are getting harder to find after initial failure (decreasing $\nu(\tau)$).

Examples of Thinking Value, IV

An arrival on the thinking arm solves the problem but the solution materializes only after a random delay.

Solution is tangible after time Δ which is exponentially distributed with rate ν .

A success on this arm has value B_ν .

Pulling this arm has a flow cost of c_ν .

The delay captures the idea that

- data from customer research has to be analyzed sufficiently.
- a convincing pitch of the pivot has to be prepared.

Examples of Thinking Value, V

An arrival on the thinking arm triggers a profit stream $db(t)$ until the deadline.

Flow profits from launching in some market.

Flow profits follow an Ornstein-Uhlenbeck process with

$$db(t) = \nu(B_\nu - b(t))dt + \sigma dW_t, \quad b(0) = 0.$$

B_ν is the long-run expected profit.

ν the rate of mean reversion.

Abandoning the Doing Arm

Assume for simplicity that an arrival on the thinking arm leads to abandoning the doing arm for good.

Can give microfoundation such that abandoning is endogenously optimal.

1. New arm replaces the old arm.
2. Belief about/payoff of new arm is high.
3. Cost of holding an arm idle (Forand, 2015).
4. Cost of switching between implementation arms.

Implications

Implications: Entrepreneurial Problem Solving

For simplicity, consider Example 1.

- Arrival on the thinking arm triggers a new exponential bandit arm with known arrival rate ν , same payoff B and flow cost c as doing arm.



Proposition

Frontloading doing reduces the probability of solving the problem.

Whenever there is an initial doing phase, the agent pulls the doing arm initially to reduce the expected time worked. This reduces the success probability as thinking successes arrive with less time remaining.

The action bias ($c > 0$) leads to false starts (Eisenmann, 2021) whenever the deadline—given the initial belief—is not too close.

False Starts



False Starts and Deadlines

Proposition

The length of the initial doing period, τ_1 , and the length of the thinking period are nondecreasing in the deadline T .

False starts become more serious as the deadline increases.

But false starts become second order as $\lim_{T \rightarrow \infty} \tau_2(T) = \infty$.

Time Pressure Can Encourage Thinking Solutions

False starts occur because entrepreneurs delay customer research. Move forward to reach a milestone quickly without initial effort to identify a potentially better route to success.

When false starts are a problem, they occur because:


1. Doing can yield a quick success.
2. Time pressure is low enough such that thinking later is still attractive.

Shorter deadlines increase time pressure; thus, discourage early thinking.

An increase in time pressure can increase the probability that the right market is identified via customer research.

A venture capitalist can use the deadline to incentivize the entrepreneur to take the desired road.

Higher Beliefs Lower Success Probability



Higher initial beliefs encourage false starts and can lead to a lower overall success probability if deadline is long enough.

Discussion of Assumptions

Relative Concavity

The main assumption driving our results is the relative concavity assumption

$$-\frac{V''(\tau)}{V'(\tau)} \geq \bar{p}\lambda.$$

However, the assumption is only sufficient to yield our results but easy to check (and satisfy).

Ensures that: deadline effect reduces faster ($V''(\tau)$) in the remaining than the payoff-on-arrival effect ($p_\tau \lambda V'(\tau)$). Time pressure is driving force behind agent's choices.

Without this assumption: if the thinking arm can have a payoff advantage over the doing arm, the agent may think twice. ▶ No Relative Concavity

Value of Thinking Independent of A_τ

Assumption so far: belief about doing arm irrelevant for value of arrival on thinking arm.

Might not be realistic:

- Agent may return to doing arm when pessimistic about new method.
- Arms may be correlated.

Structure of optimal policy extends under adjusted relative concavity assumption.

► Generalized Result

For example, when agent can return to doing arm and mixes thereafter. ► Example

Conclusion

Conclusion

Problem solving under time pressure when agent faces a time-risk tradeoff.

- Doing is fast but fundamentally risky.
- Thinking can be safe but might be too slow.

The agent never thinks twice, always does before the deadline, and may start by doing.

Rationalize one frequent reason of startup failure: false starts (Eisenmann, 2021).

- Entrepreneurs have a tendency to do early and thereby reduce their overall chances of success.

Applications

Model applicable beyond entrepreneurial problem solving.

Investment patterns for development of nuclear fusion reminiscent of optimal policy.

Politicians may have an action bias to address problems quickly rather than think about fundamental solutions.

- Paris agreement—focus on minor policies rather than much needed transformational change?

Shorter tenure clocks may lead to more positive tenure decisions—even with same requirements.

- Action bias: With longer deadline, try Top 5 with job market paper before submitting to Top Field.

Optimal Control Problem

$$\mathcal{H} = e^{-\mu(T-\tau-A_\tau)}(1 - \bar{p})\left((1 - a_\tau)\mu V(\tau) - c\right) \\ + e^{-\mu(T-\tau-\mu A_\tau)}\bar{p}e^{-\lambda A_\tau}\left((1 - a_\tau)\mu V(\tau) + a_\tau\lambda B - c\right) + a_\tau\eta_\tau$$

s. t. $\dot{A}_\tau = -a_\tau$

$$\dot{\eta}_\tau = e^{-\mu(T-\tau-A_\tau)}\left(\mu(1 - \bar{p})\left((1 - a_\tau)\mu V(\tau) - c\right) - (\lambda - \mu)e^{-\lambda A_\tau}\bar{p}\left((1 - a_\tau)\mu V(\tau) + a_\tau\lambda B - c\right)\right)$$

Technical Condition

If $\mu > \lambda$, then

$$\frac{\mu V''(\tau)}{(\lambda - \mu)\lambda^2 e^{-\lambda\tau} (B - c/\lambda)} > \frac{\mu (V(\tau) + c\tau)}{\mu(B + c\tau) + (\lambda - \mu) (B - (1 - e^{-\lambda\tau}) (B - \frac{c}{\lambda}))}.$$

Ensures that the agent is indifferent between doing for the time remaining and thinking for an instant before doing for the time remaining at most once. Again, is only a sufficient condition. [◀ back](#)

Algorithm - Preliminaries, I

We need three elements to set up our algorithm.

$$q(\tau) := \min \left\{ 1, \frac{\mu(V(\tau) + c\tau)}{\mu(B + c\tau) + (\lambda - \mu)(B - (1 - e^{-\lambda\tau})(B - \frac{c}{\lambda}))} \right\}$$

$q(\tau)$ defines the belief at which the agent is indifferent between doing for the remaining time τ and spending one more instant on the thinking arm before switching to the doing arm.

$$\dot{y}(s; p, \xi) := \left. \frac{dy_{\xi+s}}{ds} \right|_{y_{\xi}=0} = \mu V'(\xi + s) + p\mu\lambda(V(s + \xi) - B) + (\mu - \lambda p)c, \text{ and}$$

$$\hat{y}(\tau; p, \xi) := \int_0^{\tau} e^{\mu s} \dot{y}(s; p, \xi) ds.$$

\hat{y} finds root of switching function defining length of thinking period.

Algorithm - Preliminaries, II

Given any τ_3 , we can find

$$\tau_1(\tau_3) := \frac{1}{\lambda} \ln \left(\frac{\bar{p}}{1 - \bar{p}} \frac{1 - q(\tau_3)}{q(\tau_3)} \right), \text{ and}$$

$$\tau_2(\tau_3) := \begin{cases} \min \tau > 0 \text{ s.t. } \hat{y}(\tau; q(\tau_3), \tau_3) = 0, & \text{if a root for } y \text{ given } \tau_3 \text{ exists,} \\ \infty & \text{otherwise.} \end{cases}$$

1. Set $\tau_1 = \tau_2 = \tau_3 = 0$.
2. Find the largest $\bar{\tau}_3$ such that

$$\forall t \in [0, \bar{\tau}_3] \quad q(\bar{\tau}_3 - t) \leq \frac{\bar{p}e^{-\lambda t}}{(\bar{p}e^{-\lambda t} + 1 - \bar{p})}.$$

If $\bar{\tau}_3 \geq T$, set $\tau_3 = T$, $\tau_2 = \tau_1 = 0$ and stop.

3. If $q(\bar{\tau}_3) \neq \bar{p}$ go to 5.
4. If $\tau_2(\bar{\tau}_3) \geq T - \bar{\tau}_3$, set $\tau_3 = \bar{\tau}_3$ and $\tau_2 = T - \bar{\tau}_3$ and stop.
5. Replace $\bar{\tau}_3$ by the largest $\bar{\tau}_3$ such that

$$\forall t \in [0, \bar{\tau}_3] \quad q(\bar{\tau}_3 - t) \leq \frac{q(\bar{\tau}_3)e^{-\lambda t}}{q(\bar{\tau}_3)e^{-\lambda t} + 1 - q(\bar{\tau}_3)}.$$

6. Set $\tau_3 = \bar{\tau}_3$, $\tau_1 = \tau_1(\tau_3)$ and $\tau_2 = \tau_2(\tau_3)$. If $\tau_1(\tau_3) + \tau_2(\tau_3) + \tau_3 = T$, stop. Otherwise, reduce $\bar{\tau}_3$ marginally and repeat 6.

We can still show by simply adjusting the Hamiltonian that the agent will think at most once. The required adjusted condition on $V(\tau, A_\tau)$ is

$$-\frac{\frac{d^2}{d\tau^2}V(\tau, A_\tau)}{\frac{d}{d\tau}V(\tau, A_\tau)} \geq p_\tau \lambda.$$

Condition ensures that deadline effect still dominates payoff-on-arrival and cost-to-arrival effect.

An arrival on the thinking arm generates a new exponential bandit arm which can solve the problem with \underline{p} at rate λ , generates payoff B and has flow cost c .

If the agent can return to the doing arm, the value of progress depends on A_τ as the continuation value depends on the belief about the doing arm.

If $\lambda = 1$, $\bar{p} < 2/3$, and $\underline{p} > 2/3$, then the agent's optimal policy has the same structure as in the benchmark.

No Relative Concavity

An arrival on the doing arm generates a new exponential bandit arm with known arrival rate ν with $\nu < \bar{p}\lambda$ and $B_\nu = B + c/\mu$, $c_\nu = 0$.

In this case, we can have an initial additional thinking period. Once the deadline has become sufficiently irrelevant on the thinking arm, the thinking arm has a *payoff advantage* over the doing arm. Thus, the agent may think initially.

