

Knowing when to stop

Probability: the science of chance

Adriana Ocejo

Department of Statistics
University of Warwick

Postgraduate Research Lunches
4th March 2013

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

- Introduction
- Optimal stopping problems
- A simple example
- Research
- Final words

Every decision is a risky business

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

- In industry: best time to launch a new model, to sell a house, to stop trading on the financial market, etc.
- In life: best time to book a flight, to accept a job offer when waiting for other responses, etc.



No one predicts the future with full certainty... but

Probability may improve the odds of MAKING A GOOD CHOICE.

Every decision is a risky business

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

- In industry: best time to launch a new model, to sell a house, to stop trading on the financial market, etc.
- In life: best time to book a flight, to accept a job offer when waiting for other responses, etc.



No one predicts the future with full certainty... but

Probability may improve the odds of **MAKING A GOOD CHOICE**.

Every decision is a risky business

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

- In industry: best time to launch a new model, to sell a house, to stop trading on the financial market, etc.
- In life: best time to book a flight, to accept a job offer when waiting for other responses, etc.



No one predicts the future with full certainty... but

Probability may improve the odds of **MAKING A GOOD CHOICE**.

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

Can we stop a random process in a way that we optimize our objective?

A decision maker observes a random process X evolving in time.

At time t , and based on what is known about X up to that time, he or she must make a decision: either stop or continue the process.

The aim is to optimize our objective, namely

- to maximize a reward, or
- to minimize a cost.

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

Can we stop a random process in a way that we optimize our objective?

A decision maker observes a random process X evolving in time.

At time t , and based on what is known about X up to that time, he or she must make a decision: either stop or continue the process.

The aim is to optimize our objective, namely

- to maximize a reward, or
- to minimize a cost.

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

Can we stop a random process in a way that we optimize our objective?

A decision maker observes a random process X evolving in time.

At time t , and based on what is known about X up to that time, he or she must make a decision: either **stop** or **continue** the process.

The aim is to optimize our objective, namely

- to maximize a reward, or
- to minimize a cost.

Can we stop a random process in a way that we optimize our objective?

A decision maker observes a random process X evolving in time.

At time t , and based on what is known about X up to that time, he or she must make a decision: either **stop** or **continue** the process.

The aim is to optimize our objective, namely

- to maximize a reward, or
- to minimize a cost.

- If f is the reward function and the process is stopped at time t , the gain AT THAT time is $f(X_t)$.
- Since X is random, the EXPECTED reward is $\mathbb{E} f(X_t)$.
- A *stopping time* τ is a rule or strategy to stop the process (and it's also random).
- Then we want to solve the Optimal Stopping Problem:

$$\text{maximize } \mathbb{E} f(X_\tau),$$

over all the possible stopping times τ .

- If f is the reward function and the process is stopped at time t , the gain AT THAT time is $f(X_t)$.
- Since X is random, the EXPECTED reward is $\mathbb{E} f(X_t)$.
- A *stopping time* τ is a rule or strategy to stop the process (and it's also random).
- Then we want to solve the Optimal Stopping Problem:

$$\text{maximize } \mathbb{E} f(X_\tau),$$

over all the possible stopping times τ .

- If f is the reward function and the process is stopped at time t , the gain AT THAT time is $f(X_t)$.
- Since X is random, the EXPECTED reward is $\mathbb{E} f(X_t)$.
- A *stopping time* τ is a rule or strategy to stop the process (and it's also random).
- Then we want to solve the Optimal Stopping Problem:

$$\text{maximize } \mathbb{E} f(X_\tau),$$

over all the possible stopping times τ .

- If f is the reward function and the process is stopped at time t , the gain AT THAT time is $f(X_t)$.
- Since X is random, the EXPECTED reward is $\mathbb{E} f(X_t)$.
- A *stopping time* τ is a rule or strategy to stop the process (and it's also random).
- Then we want to solve the Optimal Stopping Problem:

$$\text{maximize } \mathbb{E} f(X_\tau),$$

over all the possible stopping times τ .

Example: best choice problem

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

Suppose that we want to hire the *best* secretary out of 100 candidates and that

- The interviews are arranged randomly.
- We have no information about candidates we haven't yet spoken to.
- We are able to rank each candidate interviewing relative to the ones seen by that point.
- After each interview a decision has to be made, either we *stop* and hire that person, or *continue* and reject that person FOREVER to interview the next candidate.

Intuition: certainly we can simply stop in the first interview, in that case our chance of choosing the best candidate is 1 out of 100...

CAN WE DO BETTER?

CAN WE MAXIMIZE OUR CHANCE TO HIRE THE BEST?

Yes, there is a strategy that will increase our chance to 37%!

Example: best choice problem

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

Suppose that we want to hire the *best* secretary out of 100 candidates and that

- The interviews are arranged randomly.
- We have no information about candidates we haven't yet spoken to.
- We are able to rank each candidate interviewing relative to the ones seen by that point.
- After each interview a decision has to be made, either we *stop* and hire that person, or *continue* and reject that person FOREVER to interview the next candidate.

Intuition: certainly we can simply stop in the first interview, in that case our chance of choosing the best candidate is 1 out of 100...

CAN WE DO BETTER?

CAN WE MAXIMIZE OUR CHANCE TO HIRE THE BEST?

Yes, there is a strategy that will increase our chance to 37%!

Example: best choice problem

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

Suppose that we want to hire the *best* secretary out of 100 candidates and that

- The interviews are arranged randomly.
- We have no information about candidates we haven't yet spoken to.
- We are able to rank each candidate interviewing relative to the ones seen by that point.
- After each interview a decision has to be made, either we *stop* and hire that person, or *continue* and reject that person FOREVER to interview the next candidate.

Intuition: certainly we can simply stop in the first interview, in that case our chance of choosing the best candidate is 1 out of 100...

CAN WE DO BETTER?

CAN WE MAXIMIZE OUR CHANCE TO HIRE THE BEST?

Yes, there is a strategy that will increase our chance to 37%!

Example: best choice problem

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

Suppose that we want to hire the *best* secretary out of 100 candidates and that

- The interviews are arranged randomly.
- We have no information about candidates we haven't yet spoken to.
- We are able to rank each candidate interviewing relative to the ones seen by that point.
- After each interview a decision has to be made, either we *stop* and hire that person, or *continue* and reject that person FOREVER to interview the next candidate.

Intuition: certainly we can simply stop in the first interview, in that case our chance of choosing the best candidate is 1 out of 100...

CAN WE DO BETTER?

CAN WE MAXIMIZE OUR CHANCE TO HIRE THE BEST?

Yes, there is a strategy that will increase our chance to 37%!

Solution: $1/e$ law

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

The number e is Euler's number and $1/e \approx .37$



The strategy that maximizes the probability of selecting the best candidate is as follows:

- 1 Call a candidate a *record* if he or she is the best candidate interviewed so far.
- 2 Interview and reject the first 37 candidates.
- 3 Accept (stop) the next record or select the last otherwise.

Your chance of selecting the best candidate is about 1 out of 3!

Solution: $1/e$ law

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

The number e is Euler's number and $1/e \approx .37$



The strategy that maximizes the probability of selecting the best candidate is as follows:

- 1 Call a candidate a *record* if he or she is the best candidate interviewed so far.
- 2 Interview and reject the first 37 candidates.
- 3 Accept (stop) the next record or select the last otherwise.

Your chance of selecting the best candidate is about 1 out of 3!

Solution: $1/e$ law

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

The number e is Euler's number and $1/e \approx .37$



The strategy that maximizes the probability of selecting the best candidate is as follows:

- 1 Call a candidate a *record* if he or she is the best candidate interviewed so far.
- 2 Interview and reject the first 37 candidates.
- 3 Accept (stop) the next record or select the last otherwise.

Your chance of selecting the best candidate is about 1 out of 3!

Solution: $1/e$ law

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

The number e is Euler's number and $1/e \approx .37$



The strategy that maximizes the probability of selecting the best candidate is as follows:

- 1 Call a candidate a *record* if he or she is the best candidate interviewed so far.
- 2 Interview and reject the first 37 candidates.
- 3 Accept (stop) the next record or select the last otherwise.

Your chance of selecting the best candidate is about 1 out of 3!

Variations of the problem

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

The best-choice problem is also known as

- The marriage problem: to select your **life partner** from a certain number of candidates (by Merrill Flood in 1949).
- The game of googol: to find the **biggest number** out of "as many as you wish with numbers ranging from 1 to the size of a googol".



GOOGOL
10¹⁰⁰

The strategy applies and does not depend on the number of candidates:

37% chance of selecting the best partner, number, etc!

Variations of the problem

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

The best-choice problem is also known as

- The marriage problem: to select your **life partner** from a certain number of candidates (by Merrill Flood in 1949).
- The game of googol: to find the **biggest number** out of "as many as you wish with numbers ranging from 1 to the size of a googol".



GOOGOL
10¹⁰⁰

The strategy applies and does not depend on the number of candidates:

37% chance of selecting the best partner, number, etc!

I have studied Optimal Stopping Problems with applications in Finance, to answer questions like:

When it is optimal to stop trading on financial markets to maximize our reward, for certain contracts?

What analytical properties does the optimized value possess and under which conditions?

I currently study Optimal Stopping and Control Problems, e.g.

- Besides a stopping rule we also choose a parameter that controls the process X , or
- There are two players, one chooses the stopping rule to maximize whereas the other one chooses the control to minimize (as a game).

I have studied Optimal Stopping Problems with applications in Finance, to answer questions like:

When it is optimal to stop trading on financial markets to maximize our reward, for certain contracts?

What analytical properties does the optimized value possess and under which conditions?

I currently study Optimal Stopping and Control Problems, e.g.

- Besides a stopping rule we also choose a parameter that controls the process X , or
- There are two players, one chooses the stopping rule to maximize whereas the other one chooses the control to minimize (as a game).

I have studied Optimal Stopping Problems with applications in Finance, to answer questions like:

When it is optimal to stop trading on financial markets to maximize our reward, for certain contracts?

What analytical properties does the optimized value possess and under which conditions?

I currently study Optimal Stopping and Control Problems, e.g.

- Besides a stopping rule we also choose a parameter that controls the process X , or
- There are two players, one chooses the stopping rule to maximize whereas the other one chooses the control to minimize (as a game).

I have studied Optimal Stopping Problems with applications in Finance, to answer questions like:

When it is optimal to stop trading on financial markets to maximize our reward, for certain contracts?

What analytical properties does the optimized value possess and under which conditions?

I currently study Optimal Stopping and Control Problems, e.g.

- Besides a stopping rule we also choose a parameter that controls the process X , or
- There are two players, one chooses the stopping rule to maximize whereas the other one chooses the control to minimize (as a game).

Applications of OSP in other areas different of Finance

Knowing when to stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

- In Analytic Philosophy [5], *to improve the performance of ontology refinement methods.*
- In Operations Management [4], *to determine the time to introduce a new product to the market.*
- In Biology [1], *in modeling biodiversity loss problems (to determine the optimal point to stop the conversion process of replacing naturally existing species by human chosen ones).*

Something I've learnt during my PhD journey

Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping

Example

Research

Final words

References

*Patience is not about how LONG you can wait...
It's about how WELL you behave while you wait.*



Knowing when to
stop

Adriana Ocejo

Overview

Introduction

Optimal stopping






Example

Research

Final words

References



-  [1] AMITRAJEET, A., *An Optimal Stopping approach to the Conservation of Biodiversity*, Economic Research Institute Study Paper, Utah State University, (1997).
-  [2] BRUSS, T., *Sum the Odds to One and Stop*, The Annals of Probability, Vol. 28, 1384-1391, (2000).
-  [3] HILL, T. *Knowing When to Stop*, American Scientist, Vol 97, 126-133, (2009).
-  [4] SECHAN, O., *Optimal Stopping Problems in Operations Management*, Dissertation for the degree of Doctor of Philosophy, Stanford University, (2010).
-  [5] WEICHSELBRAUN, A., *Applying Optimal Stopping Theory to Improve the Performance of Ontology Refinement Methods*, System Sciences (HICSS) Conference Publication, (2011).