Time-Consistent Stopping

Yu-Jui Huang
Dublin City University

Joint work with
Adrien Nguyen-Huu
École Polytechnique

Boston University
February 1, 2016

OUTLINE

Motivation

- What is time inconsistency? Why do we have it?

Methodology

- Game-theoretic approach

Main Results

Extensions

Classical Optimal Stopping

Consider

- a continuous Markovian process $X : [0, \infty) \times \Omega \mapsto \mathbb{R}^d$.
- a continuous payoff function $g : \mathbb{R}^d \mapsto \mathbb{R}_+$.

Optimal Stopping

Given $(t, x) \in [0, \infty) \times \mathbb{R}^d$, can we solve

$$\sup_{\tau \in \mathcal{T}_t} \mathbb{E}_{t,x}[\delta(\tau - t)g(X_\tau)]?$$

- $\mathcal{T}_t$: set of stopping times τ s.t. $\tau \geq t$ a.s.
- $\delta : \mathbb{R}_+ \mapsto [0, 1]$: decreasing from $\delta(0) = 1$.

Optimal Stopping Times [Karatzas & Shreve (1998)]

For all $(t,x) \in \mathbb{R}_+ \times \mathbb{R}^d$, the stopping time

$$\widetilde{\tau}(t,x) := \inf \Big\{ s \geq t : \delta(s-t) g(X_s^{t,x}) = \operatorname*{ess\,sup}_{\tau \in \mathcal{T}_s} \mathbb{E}_{s, X_s^{t,x}}[\delta(\tau - t) g(X_\tau)] \Big\}$$

is optimal, i.e.

$$\mathbb{E}_{t,x}[\delta(\widetilde{\tau}(t,x) - t) g(X^{t,x}_{\widetilde{\tau}(t,x)})] = \sup_{\tau \in \mathcal{T}_t} \mathbb{E}_{t,x}[\delta(\tau - t) g(X_\tau)].$$

We say $\widetilde{\tau}$ is a **stopping policy**:

$$(t,x) \longmapsto \widetilde{\tau}(t,x) \in \mathcal{T}_t$$

Classical Theory: END OF STORY!

- **Problem Solved.** *Feeling Good?*

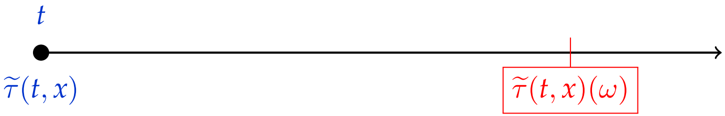

- **The Reality:**

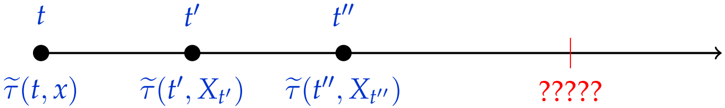

- **Time Inconsistency:**
  - $\widetilde{\tau}(t,x)$, $\widetilde{\tau}(t',X_{t'})$, $\widetilde{\tau}(t'',X_{t''})$ may all be different.
  - Is it reasonable to apply $\widetilde{\tau}(t,x)$ at time t?

EXAMPLE (BES 1)

- X_t : one-dimensional Brownian motion
- Hyperbolic discount function

$$\delta(s) = \frac{1}{1+s}.$$

- payoff function $g(x) = |x|$.

Using PDE approach, we solve explicitly

$$\tilde{\tau}(t,x) = \inf\left\{s \geq t : |X_s^{t,x}| \geq \sqrt{1+(s-t)}\right\}.$$

- Free boundary $s \mapsto \sqrt{1+(s-t)}$ depends on initial time t.
- This induces **time inconsistency**.

Free boundary $s \mapsto \sqrt{1+(s-t)}$ is changing over time t.

- it keeps moving to the right.

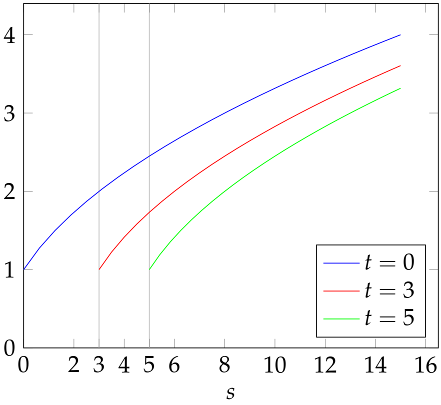

- $\widetilde{\tau}(t,x) = \inf\{s \geq t : |X_s^{t,x}| \geq \sqrt{1+(s-t)}\}$ inconsistent over time.

Safe Case: Exponential Discounting

In classical literature of Mathematical Finance,

$$\delta(s) = e^{-\rho \cdot s} \quad \text{for some } \rho > 0.$$

- This means $\boxed{\delta(s-t)\delta(r-s) = \delta(r-t)}, \forall\, 0 \le t \le s \le r.$
- Optimal stopping time becomes

$$\widetilde{\tau}(t,x) := \inf \left\{ s \ge t \;:\; \delta(s-t) g(X_s) = \operatorname*{ess\,sup}_{\tau \in \mathcal{T}_s} \mathbb{E}_{s,X_s}[\delta(\tau - t) g(X_\tau)] \right\}$$

$$= \inf \left\{ s \ge t \;:\; g(X_s) = \operatorname*{ess\,sup}_{\tau \in \mathcal{T}_s} \mathbb{E}_{s,X_s}[\delta(\tau - s) g(X_\tau)] \right\}.$$

No t-dependence anymore!

Why not stay with exponential discounting?

- Payoff may not be monetary (utility, happiness, health,...).
- **Empirical:** people don't discount money exponentially.
  - People admit "decreasing impatience" (Laibson (1997), O'Donoghue & Rabin (1999))

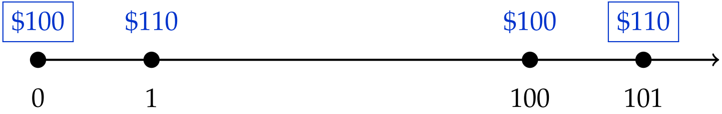

$\frac{100}{\delta(1)200} > 1$ $\qquad$ $\frac{\delta(100)100}{\delta(101)200} < 1$

  - If $\delta(s-t) = e^{-\rho(s-t)}$,

$$\frac{100}{\delta(1)200} = \frac{\delta(100)100}{\delta(101)200} = \frac{e^{\rho}}{2} \text{ is constant.}$$

$\Rightarrow$ Does not capture "decreasing impatience".

Literature

Stroz (1955): 3 different reactions to time inconsistency

- A **naive agent** follows classical optimal stopping.
- A **pre-committed agent** forces all his future selves to follow the initial optimal stopping time $\widetilde{\tau}(t, x)$.
- A **sophisticated** **agent**
  1. considers the behavior of future selves;
  2. aims to find a stopping strategy that

once being enforced over time,
no future self would want to deviate from it.

Question: How to formulate sophisticated strategies in **continuous time** ?

Unclear in the literature...

LITERATURE

- Ekeland & Lazrak (2006): **Subgame perfect Nash equilibriums** emerge as the proper formulation for sophisticated strategies, for **control problems**.

$$\textbf{sophisticated strategies} \iff \textbf{equilibrium strategies}$$

- Recent studies: Ekeland & Pirvu (2008), Ekeland, Mbodji, & Pirvu (2012), Björk, Murgoci, & Zhou (2014), Dong & Sircar (2014), Björk & Murgoci (2014), Yong (2012),...

- Extending the equilibrium idea to **stopping problems**:

difficult, **unresolved**.

Xu & Zhou (2013), Barberis (2002), Grenadier & Wang (2007).

FORMULATING EQUILIBRIUMS

An **equilibrium strategy** is a strategy that

once being enforced over time,
no future self would want to deviate from it.

- Imagine that
  1. You select a stopping policy τ at time 0, and enforce it over time (Recall: $(t,x) \mapsto \tau(t,x) \in \mathcal{T}_t$).
  2. At time $t \geq 0$,

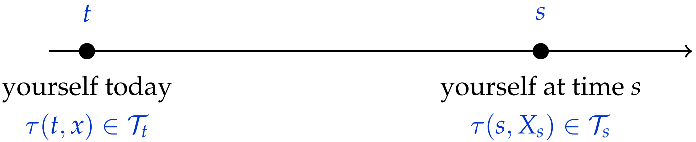

You think: Given that all future selves will use $\tau(s, X_s^{t,x})$, what is the best stopping strategy at time t?

- You feel GOOD, if $\tau(t,x)$ is the best strategy.
- You feel BAD, if $\tau(t,x)$ is not.

- **Equilibrium strategy:** a stopping policy τ s.t. when τ is enforced, all future selves feel GOOD.

Mathematical Formulation

- Pick a stopping policy τ (Recall: $(t, x) \mapsto \tau(t, x) \in \mathcal{T}_t$).

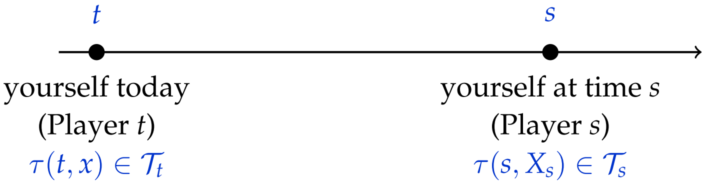

- When do we eventually stop?

$$\mathcal{L}\tau(t, x) := \inf \left\{ s \geq t \ : \ \tau(s, X_s^{t,x}) = s \right\}.$$

- **Game-theoretic** thinking of Player t:

  Given that each Player s will employ $\tau(s, X_s^{t,x}) \in \mathcal{T}_s$,
  - what is the best stopping strategy at time t?
  - can it just be $\tau(t, x)$?

BEST STOPPING STRATEGY

Player t has only **two** possible actions: to stop or to continue.

- If she stops, she gets $\boxed{g(x)}$ right away.
- If she continues, she will eventually stop at the moment

$$\mathcal{L}^*\tau(t,x) := \inf\left\{s > t \ : \ \tau(s, X_s^{t,x}) = s\right\}$$

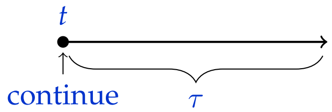

Her expected gain is therefore

$$\boxed{\mathbb{E}_{t,x}\left[\delta(\mathcal{L}^*\tau(t,x) - t)g\left(X_{\mathcal{L}^*\tau(t,x)}\right)\right]}.$$

The best stopping strategy for Player t:

I. $g(x) > \mathbb{E}_{t,x}\left[\delta(t, \mathcal{L}^*\tau(t,x))g\left(X_{\mathcal{L}^*\tau(t,x)}\right)\right] \Rightarrow$ **stop** right away

II. $g(x) < \mathbb{E}_{t,x}\left[\delta(t, \mathcal{L}^*\tau(t,x))g\left(X_{\mathcal{L}^*\tau(t,x)}\right)\right] \Rightarrow$ **continue**

- she will eventually stop at the moment $\mathcal{L}^*\tau(t,x)$.

III. $g(x) = \mathbb{E}_{t,x}\left[\delta(t, \mathcal{L}^*\tau(t,x))g\left(X_{\mathcal{L}^*\tau(t,x)}\right)\right] \Rightarrow$

- **indifferent** between to stop and to continue at time t.
- no incentive to deviate from $\tau(t,x)$
- She will eventually stop at the moment $\mathcal{L}\tau(t,x)$.

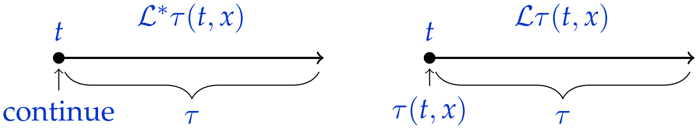

- Summarize the best stopping strategy for Player t as

$$\Theta\tau(t,x) := t\,1_{S_\tau}(t,x) + \mathcal{L}\tau(t,x)1_{I_\tau}(t,x) + \mathcal{L}^*\tau(t,x)1_{C_\tau}(t,x),$$

where

$$S_\tau := \{(t,x) : g(x) > \mathbb{E}_{t,x}\left[\delta(t,\mathcal{L}^*\tau(t,x))g\left(X_{\mathcal{L}^*\tau(t,x)}\right)\right]\},$$
$$I_\tau := \{(t,x) : g(x) = \mathbb{E}_{t,x}\left[\delta(t,\mathcal{L}^*\tau(t,x))g\left(X_{\mathcal{L}^*\tau(t,x)}\right)\right]\},$$
$$C_\tau := \{(t,x) : g(x) < \mathbb{E}_{t,x}\left[\delta(t,\mathcal{L}^*\tau(t,x))g\left(X_{\mathcal{L}^*\tau(t,x)}\right)\right]\}.$$

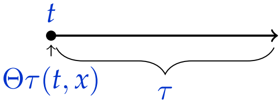

- Player t feels good to use $\tau(t,x) \iff \tau(t,x) = \Theta\tau(t,x)$.
- **Conclusions:**

All players feel good to follow τ $\iff$ $\boxed{\tau(t,x) = \Theta\tau(t,x),\ \forall(t,x)}$

EQUILIBRIUM POLICIES

Definition

A stopping policy τ is called an **equilibrium policy** if

$$\Theta\tau(t,x) = \tau(t,x) \text{ a.s.}, \quad \forall (t,x) \in [0,\infty) \times \mathbb{R}^d.$$

- **Trivial Equilibrium:** consider $\tau(t,x) := t$ for all (t,x).

$$\begin{aligned}\Theta\tau(t,x) &:= t\, 1_{S_\tau}(t,x) + \mathcal{L}\tau(t,x) 1_{I_\tau}(t,x) + \mathcal{L}^*\tau(t,x) 1_{C_\tau}(t,x)\\ &= t\, 1_{S_\tau}(t,x) + t 1_{I_\tau}(t,x) + t 1_{C_\tau}(t,x) = t = \tau(t,x).\end{aligned}$$

- **In general**, given a stopping policy τ, carry out iteration:

$$\tau \longrightarrow \Theta\tau \longrightarrow \Theta^2\tau \longrightarrow \cdots \longrightarrow \text{“equilibrium”??}$$

- **To show:** (i) $\tau_0 := \lim_{n\to\infty} \Theta^n \tau$ converges (ii) $\Theta\tau_0 = \tau_0$.

DECREASING IMPATIENCE

- **Assumption:** the discount function $\delta : \mathbb{R}_+ \mapsto [0,1]$ satisfies

$$\delta(t)\delta(s) \leq \delta(t+s) \quad \forall\, t,s \geq 0. \tag{1}$$

Definition

A discount function δ induces **Decreasing Impatience** if,

for any $s \geq 0$, $\frac{\delta(t+s)}{\delta(t)}$ is increasing in t.

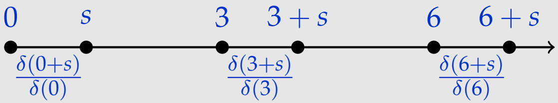

$$\text{DI} \implies \frac{\delta(t+s)}{\delta(t)} \geq \frac{\delta(0+s)}{\delta(0)} = \delta(s) \implies \delta(t)\delta(s) \leq \delta(t+s).$$

- Once we consider **DI**, (1) is automatically satisfied.

Main Result

Lemma

Assume (1). *Let* τ *be a stopping policy. Then,*

$$\text{if} \quad \boxed{\Theta\tau(t,x) \leq \tau(t,x) \;\; a.s. \quad \forall (t,x)}, \tag{2}$$

$$\text{then} \quad \Theta^{n+1}\tau(t,x) \leq \Theta^{n}\tau(t,x) \;\; a.s. \quad \forall (t,x) \text{ and } n.$$

Theorem

Assume (1) *and* (2). *Then, for any* (t,x),

$$\tau_0(t,x) := \downarrow \lim_{n\to\infty} \Theta^{n}\tau(t,x) \;\; \text{converges } a.s.$$

Moreover, τ_0 *is an equilibrium policy, i.e.*

$$\Theta\tau_0(t,x) = \tau_0(t,x) \;\; a.s. \quad \forall (t,x).$$

Recall the *classical optimal stopping time* $\widetilde{\tau}(t,x)$ for all (t,x).

- It can be shown that

$$\Theta\widetilde{\tau}(t,x) \leq \widetilde{\tau}(t,x) \text{ a.s.} \quad \text{for all } (t,x).$$

- Hence,

$$\tau_0(t,x) := \downarrow \lim_{n\to\infty} \Theta^n \widetilde{\tau}(t,x)$$

  is an equilibrium policy.

This provides a nice **economic interpretation**...

Improving via Iteration

1. At first, we follow $\widetilde{\tau}$. By game-theoretic thinking,

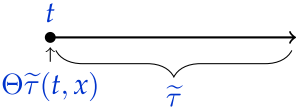

$\Longrightarrow$ switch from $\widetilde{\tau}$ to $\Theta\widetilde{\tau}$

2. Now, we follow $\Theta\widetilde{\tau}$. By game-theoretic thinking,

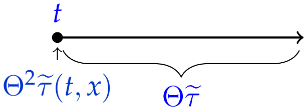

$\Longrightarrow$ switch from $\Theta\widetilde{\tau}$ to $\Theta^2\widetilde{\tau}$

3. Continue this procedure *until* we reach the equilibrium

$$\tau_0(t,x) := \downarrow \lim_{n\to\infty} \Theta^n\widetilde{\tau}(t,x)$$

Then, $\Theta\tau_0(t,x) = \tau_0(t,x) \implies$ cannot improve anymore.

FROM "NAIVE" TO "SOPHOSTICATED"

$$\tau_0 = \lim_{n\to\infty} \Theta^n \widetilde{\tau}$$

reveals the connection between "**naive**" and "**sophisticated**":

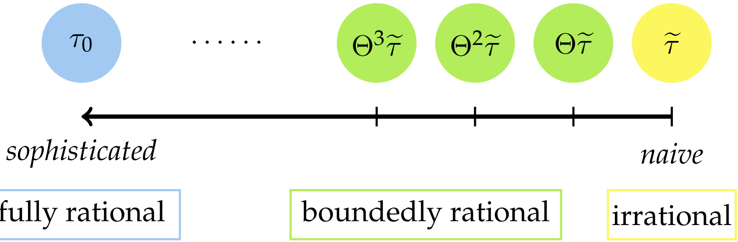

- **Bounded Rationality** proposed by H. Simon (1982).
- This connection is new in the literature.

EXAMPLE (SMOKING CESSATION)

- Smokers care most about:
  - long-term serious health problems
  - immediate pain from quitting smoking
- **Our Model:**
  - A smoker has a fixed lifetime T.
  - Deterministic cost process

$$X_s^{t,x} := xe^{\frac{1}{2}(s-t)}, \quad s \in [t, T]$$

  - Smoker can either
    - 1. quit at $s < T$ (costs X_s) 2. die peacefully at T (no cost)
    - 1. never quit (no cost) 2. die painfully at T (costs X_T)
  - Hyperbolic discounting:

$$\delta(s) = \frac{1}{1+s} \quad \forall s \geq 0.$$

- **Classical Theory:** For each $t \in [0, T]$,

$$\min_{s\in[t,T]} \delta(s-t)X_s^{t,x} = \min_{s\in[t,T]} \frac{xe^{\frac{1}{2}(s-t)}}{1+(s-t)}.$$

  - By Calculus, the optimal stopping time is

$$\widetilde{\tau}(t,x) = \begin{cases} t+1 & \text{if } t < T-1, \\ T & \text{if } t \geq T-1. \end{cases}$$

  - Observe that

$$\mathcal{L}\widetilde{\tau}(t,x) := \inf\{s \geq t \ : \ \widetilde{\tau}(s,X_s) = s\} \wedge T = T,$$
$$\mathcal{L}^*\widetilde{\tau}(t,x) := \inf\{s > t \ : \ \widetilde{\tau}(s,X_s) = s\} \wedge T = T.$$

  - **time inconsistency $\Longrightarrow$ procrastination**

- **Our Theory:** Apply equilibrium policy $\tau_0 := \lim_{n\to\infty} \Theta^n \widetilde{\tau}$.
  - First iteration:

$$\Theta\widetilde{\tau}(t,x) := t\, 1_{S_{\widetilde{\tau}}}(t,x) + \mathcal{L}\widetilde{\tau}(t,x) 1_{I_{\widetilde{\tau}}}(t,x) + \mathcal{L}^*\widetilde{\tau}(t,x) 1_{C_{\widetilde{\tau}}}(t,x),$$

$$\begin{aligned} S_{\widetilde{\tau}} &:= \{(t,x) : x < \delta(\mathcal{L}^*\widetilde{\tau}(t,x) - t) X_{\mathcal{L}^*\widetilde{\tau}(t,x)}\}, \\ I_{\widetilde{\tau}} &:= \{(t,x) : x = \delta(\mathcal{L}^*\widetilde{\tau}(t,x) - t) X_{\mathcal{L}^*\widetilde{\tau}(t,x)}\}, \\ C_{\widetilde{\tau}} &:= \{(t,x) : x > \delta(\mathcal{L}^*\widetilde{\tau}(t,x) - t) X_{\mathcal{L}^*\widetilde{\tau}(t,x)}\}. \end{aligned}$$

  - Compare $\boxed{x}$ with

$$\delta(\mathcal{L}^*\widetilde{\tau}(t,x) - t) X^{t,x}_{\mathcal{L}^*\widetilde{\tau}(t,x)} = \frac{X^{t,x}_T}{1 + (T-t)} = \boxed{x \cdot \frac{e^{\frac{1}{2}(T-t)}}{1 + (T-t)}}$$

  - Since $e^{\frac{1}{2}s} = 1 + s$ at $s = 0$ and $s^* \approx 2.513$,

$$\begin{aligned} S_{\widetilde{\tau}} &= \{(t,x) : t < T - s^*\}, \\ C_{\widetilde{\tau}} &= \{(t,x) : t \in (T - s^*, T)\}, \\ I_{\widetilde{\tau}} &= \{(t,x) : t = T - s^* \text{ or } T\}. \end{aligned}$$

- **Conclude:**

$$\Theta\widetilde{\tau}(t,x) = \begin{cases} t & \text{if } t < T - s^*, \\ T & \text{if } t \geq T - s^*. \end{cases}$$

This is already an equilibrium, i.e. $\Theta^2\widetilde{\tau} = \Theta\widetilde{\tau}$.

- Thus,

$$\tau_0(t,x) := \lim_{n\to\infty} \Theta^n\widetilde{\tau}(t,x) = \begin{cases} t & \text{if } t < T - s^*, \\ T & \text{if } t \geq T - s^*. \end{cases}.$$

- **τ_0 says "Stop Smoking Immediately!!"**
(unless you're too old...)

EXAMPLE (BES(1))

- X_t : one-dimensional Brownian motion
- Hyperbolic discount function

$$\delta(s) = \frac{1}{1+s}.$$

- payoff function $g(x) = |x|$.
- Classical optimal stopping time

$$\widetilde{\tau}(t,x) = \inf\left\{s \geq t : |X_s^{t,x}| \geq \sqrt{1+(s-t)}\right\}.$$

- Find an equilibrium policy:

$$\tau_0(t,x) := \lim_{n\to\infty} \Theta^n \widetilde{\tau}(t,x) = \Theta^3 \widetilde{\tau}(t,x) = \inf\{s \geq t : |X_s^{t,x}| \geq x^*\},$$

where x^* solves

$$\int_0^\infty e^{-s} \cosh(x\sqrt{2s}) \operatorname{sech}(\sqrt{2s}) ds = x \quad \Longrightarrow \quad x^* \approx 0.922.$$

We can characterize the whole set $\mathcal{E}$ of equilibrium policies.

- For all $a \geq 0$, define τ_a by

$$\tau_a(t,x) := \inf\{s \geq t : |X_s^{t,x}| \geq a\}, \quad \forall (t,x).$$

- $\mathcal{E} = \{\tau_a : a \in [0, a^*]\}$, where a^* solves

$$a \int_0^\infty e^{-s}\sqrt{2s}\tanh(a\sqrt{2s})ds = 1 \implies a^* \approx 0.946.$$

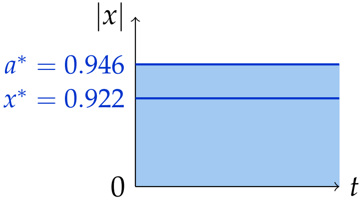

SELECTING AN EQUILIBRIUM

Question: Which equilibrium to use?

- **Optimal "time-consistent" stopping:**

$$\sup_{\tau \in \mathcal{E}} \mathbb{E}_{t,x}[\delta(\mathcal{L}\tau(t,x) - t) g(X_{\mathcal{L}\tau(t,x)})].$$

  Difficult to solve...

  - Martingale method & dynamic programming break down!
  - Know too little about $\mathcal{E}$...

- **Pareto efficiency:**
  How to formulate this under current setting?

PROBABILITY DISTORTION

- **Optimal stopping** under Probability Distortion:

$$\sup_{\tau \in \mathcal{T}_t} \int_0^\infty w\Big(\mathbb{P}_{t,x}\left[g(X_\tau) > u \right] \Big) du.$$

[Xu & Zhou (2013)]

  - This is a Choquet integral....

- Equilibrium policies can be defined similarly.
  - How to solve **Optimal time-consistent stopping**?

$$\sup_{\tau \in \mathcal{E}} \int_0^\infty w\Big(\mathbb{P}_{t,x}\left[g(X_\tau) > u \right] \Big) du.$$

THANK YOU!!

Preprint available @ arXiv:1502.03998

"Time-consistent stopping under decreasing impatience"