Time-Inconsistent Problems

Jiongmin Yong

Consider the following initial value problem of an ordinary differential equation:

$$\begin{cases}\dot{X}(s) = f(s, X(s), u(s)), & s \in [t, T], \\ X(t) = x. \end{cases} \tag{1}$$

In this differential equation, $u(\cdot)$, the *control*, is selected from some class $\mathcal{U}[t, T]$ of measurable functions with values in a metric space U, and the solution $X(\cdot)$ of (1) is called the *state process*. In applications, $X(s)$ could be the total wealth from some stocks, tonnes of fish in a large lake, or the location of a moving object at time s, and $u(\cdot)$ could be a trading strategy, a harvest rate, or a driving force. Different selections of controls $u(\cdot)$ lead to different state processes $X(\cdot)$. To measure the performance of the control toward a goal, one may introduce the following cost functional:

$$J(t, x; u(\cdot)) = \int_t^T e^{-\delta(s-t)} g(X(s), u(s)) ds. \tag{2}$$

In this integral, g is the *running cost rate* and $e^{-\delta(s-t)}$ is the *exponential discounting* (with *discount rate* $\delta > 0$). To find a control that minimizes the cost functional, a classical optimal control problem can be stated as follows.

Problem (C). For a given *initial pair* $(t, x) \in [0, T] \times \mathbb{R}^n$, find $\bar{u}(\cdot) \in \mathcal{U}[t, T]$ such that

$$J(t, x; \bar{u}(\cdot)) = \inf_{u(\cdot) \in \mathcal{U}[t,T]} J(t, x; u(\cdot)) \equiv V(t, x).$$

In the above, $\bar{u}(\cdot)$ is an *optimal control* and $V(\cdot\,, \cdot)$ is defined as the *value function*. Bellman's *principle of optimality* can be stated as: for any $(t, x) \in [0, T) \times \mathbb{R}^n$ and $\tau \in (t, T]$,

$$V(t, x) = \inf_{u(\cdot) \in \mathcal{U}[t,\tau]} \Big\{ \int_t^\tau e^{-\delta(s-t)} g(X(s), u(s)) ds + e^{-\delta(\tau-t)} V(\tau, X(\tau)) \Big\}.$$

This principle gives: if $\bar{u}(\cdot)$ is an optimal control at (t, x) with $\bar{X}(\cdot)$ being the corresponding optimal state process, then for any $\tau \in (t, T]$,

$$\begin{aligned} V(t, x) &= J(t, x; \bar{u}(\cdot)) \\ &= \int_t^\tau e^{-\delta(s-t)} g(\bar{X}(s), \bar{u}(s)) ds + e^{-\delta(\tau-t)} J(\tau, \bar{X}(\tau); \bar{u}(\cdot)) \\ &\geqslant \int_t^\tau e^{-\delta(s-t)} g(\bar{X}(s), \bar{u}(s)) ds + e^{-\delta(\tau-t)} V(\tau, \bar{X}(\tau)) \\ &\geqslant V(t, x). \end{aligned}$$

This implies

$$V(\tau, \bar{X}(\tau)) = J\big(\tau, \bar{X}(\tau); \bar{u}(\cdot)\big|_{[\tau,T]}\big),$$

which means the restriction $\bar{u}(\cdot)\big|_{[\tau,T]}$ of an optimal control $\bar{u}(\cdot)$ selected for the initial pair (t, x) on $[t, T]$ is an optimal control for the initial pair $(\tau, \bar{X}(\tau))$ on $[\tau, T]$. In other words, an optimal control $\bar{u}(\cdot)$ determined at time t will stay optimal thereafter. Such a phenomenon is called the *time-consistency* of the control problem or the optimal control.

People frequently regret the decision they made earlier, meaning an optimal decision made today will hardly seem optimal forever. We call such a phenomenon the *time-inconsistency* of the problem under consideration. There are two major reasons causing time-inconsistency: *time-preferences* and *risk-preferences*. To elaborate a bit more, normal people usually are not 100% rational and they often over-weigh the immediate satisfaction level (the utility) or regard the immediate time period more precious. Here is an example: if you are invited to referee a paper, would you review the paper immediately or would you wait until the associate editor sends you a reminder? Time-preferences play a role here. Different people could have different opinions on upcoming uncertain events. An easy example to illustrate risk-preferences is the opinion on whether to buy a risky stock.

Mathematically, time-preferences can be described by *discounting*, and risk-preferences can be described by *subjective probability*. It is possible to present such problems in a stochastic setting. But for simplicity, let us continue with the deterministic case. Problem (C) (with exponential discounting) is time-consistent, which represents a situation in which the controller (the person who is controlling the system) is rational. Now if the controller is not 100% rational, then the discounting might not be exponential. A typical nonexponential discounting is *hyperbolic discounting*, which is illustrated as

$$J(t, x; u(\cdot)) = \int_t^T \frac{1}{1 + a(s-t)} g(X(s), u(s)) ds \tag{3}$$

for some $a > 0$. With such a cost functional, the corresponding optimal control problem will be time-inconsistent, i.e., if $\bar{u}(\cdot)$ is an optimal control for (t, x) defined on $[t, T]$ with optimal state process $\bar{X}(\cdot)$, then there

Jiongmin Yong is a professor of mathematics at the University of Central Florida. His email address is jiongmin.yong@ucf.edu.

DOI: https://doi.org/10.1090/noti2242

will be a later time $\tau \in (t, T]$ such that the optimal control $\hat{u}(\cdot)$ for $(\tau, \bar{X}(\tau))$ defined on $[\tau, T]$ does not satisfy the following:

$$\hat{u}(s) = \bar{u}(s), \qquad \text{a.e. } s \in [\tau, T].$$

For time-inconsistent optimal control problems, it is not wise to find optimal controls. Instead, one should look for *equilibrium strategies*. The idea is to regard the problem as a multiperson differential game in which today's self plays with future selves. To make the future reasonably satisfactory, today's self should be willing to sacrifice some immediate satisfaction. Saving for retirement is such an example. Under proper conditions, such equilibrium strategies can be constructed.

Jiongmin Yong

Credits

Photo of Jiongmin Yong is courtesy of Jiongmin Yong.