R07323033 錢嘉鴻

What is the question?

The techniques used to solve stochastic games in discrete time often depend crucially on the form of the underlying state process. The analysis of irreducible games rely on convergence to a steady state distribution , whereas the analysis of absorbing games often hinges on the fact that there is only one non-absorbing state.

Why should we care?

The techniques in this paper do not rely on any such properties, hence they are applicable to stochastic games with an arbitrary state process

A real world application?

In the global financial crisis, financial institutions were unable to raise sufficient capital to meet their short-term liabilities because investors had lost confidence in the financial sector. The past performance of the financial sector had led credit lines to dry up and created an environment where financial institutions were not able to bridge short-term liquidity gaps in the usual way. The development of new technologies by competing research institutions exhibits a similar history-dependent environment. The successful discovery of a new technology changes the research environment forever: competing researchers will not be able to patent similar work anymore and any effort put into such a discovery was exerted in vain. It is impossible to forecast the exact time of a financial crisis or the discovery of a new technology. The occurrence of such a state change is random and the likelihood depends on the involved parties' actions. A game-theoretic model that accounts for these sudden state changes is a stochastic game. No deterministic-time dynamic game can capture these sudden and potentially drastic changes in the environment.

What is the author's answer?

Based on recent developments in continuous-time repeated games, this paper provides a unifying framework for the analysis of stochastic games with imperfect public monitoring in a continuous-time setting. The methodology is not limited to irreducible games or absorbing games and it is applicable to any stochastic game, as long as the public signal satisfies Assumptions 1 and 2.

Assumption 1. For each y ∈ Y, every action profile a ∈ A(y) has pairwise full rank.

Assumption 2. span M1 y (a) ⊥ span M2 y (a) for each y ∈ Y and each a ∈ A(y).

我覺得這是一套研究方法，所以解答就是....這套過程

How did the author get there? (1) Methodlogy

Definition 2.1. A *(public) strategy* A^i of player i is an $\mathbb{F}$-predictable process that takes values in $\mathcal{A}^i(S_-)$. We denote by $A = (A^1, A^2)$ a *(public) strategy profile.*

Definition 2.2.

(i) Each player i receives an unobservable expected flow payoff $g^i : \mathcal{Y} \times \mathcal{A} \to \mathbb{R}$.[6]

(ii) Player i's *discounted expected future payoff* (or *continuation value*) under strategy profile A at any time $t \geq 0$ is given by

$$W_t^i(S_t, A) = \int_t^\infty r\mathrm{e}^{-r(s-t)} \mathbb{E}_{Q_s^A}\left[g^i(S_s, A_s) \,\middle|\, \mathcal{F}_t\right] \mathrm{d}s, \tag{2}$$

where $r > 0$ is the discount rate of both players and the distribution of $(S_s)_{s \geq t}$ is determined uniquely by S_t and $(A_s)_{s \geq t}$.

(iii) A strategy profile A is a *perfect public equilibrium (PPE)* for discount rate r if for

Definition 2.3.

(i) A PPE A is a *stationary Markov (perfect) equilibrium* for initial state S_0 if there exists a map $a^* : \mathcal{Y} \to \mathcal{A}$ with $a_*(y) \in \mathcal{A}(y)$ for every state y such that $A = a^*(S_-)$.

(ii) A PPE is *semi-stationary* for initial state S_0 if there exists a map $a_* : \bigcup_{k=1}^{\infty} \mathcal{Y}^k \to \mathcal{A}$ with $a_*(y_1, \ldots, y_k) \in \mathcal{A}(y_k)$ for any sequence of states of any length k such that $A = a_*(\hat{S}_-)$, where $\hat{S}_t$ is the sequence of states visited up to and including time t.

(iii) We denote by $\mathcal{E}^M(r)$ and $\mathcal{E}^S(r)$ the families of payoff pairs that are achievable in stationary Markov and semi-stationary equilibria, respectively. Both of those are without the use of a public randomization device. We denote by $\mathcal{E}_p^S(r)$ the family of semi-stationary equilibrium payoffs with public randomization. Note that

$$\mathcal{E}^M(r) \subseteq \mathcal{E}^S(r) \subseteq \mathcal{E}_p^S(r) \subseteq \mathcal{E}(r) \subseteq \mathcal{V}^*.$$

every player i and all possible deviations $\tilde{A}^i$,

$$W^i(S, A) \geq W^i\big(S, (\tilde{A}^i, A^{-i})\big) \text{ a.e.,}^7 \tag{3}$$

where A^{-i} denotes the strategy of player i's opponent in profile A.

(iv) We denote the set of all payoff vectors that are achievable by a perfect public equilibrium when the initial state is y and the discount rate is r by

$$\mathcal{E}_y(r) := \big\{w \in \mathbb{R}^2 \bigm| \text{there exists a PPE } A \text{ with } W_0(y, A) = w \text{ a.s.}\big\}.$$

We denote by $\mathcal{E}(r)$ the family $\big(\mathcal{E}_y(r)\big)_{y \in \mathcal{Y}}$ of equilibrium payoff sets.

7 Computation

In the terminology of Markov processes, a set of states $\mathcal{Y}_0 \subseteq \mathcal{Y}$ is a *communicating class* if each state within $\mathcal{Y}_0$ can be reached from any other state in $\mathcal{Y}_0$ with positive probability (either directly or indirectly). Computation of the family of equilibrium payoff sets proceeds by communicating classes. Communicating classes of a Markov process can be organized in a directed acyclic graph as illustrated in Figure 11.[14] A communicating class $\mathcal{Y}_0$ is a *direct predecessor class* of $\mathcal{Y}_1$, denoted $\mathcal{Y}_0 \prec \mathcal{Y}_1$, if some state in $\mathcal{Y}_1$ can be reached directly from some state in $\mathcal{Y}_0$. If $\mathcal{Y}_0 \prec \mathcal{Y}_1$, we also say that $\mathcal{Y}_1$ is a *direct successor class* of $\mathcal{Y}_0$. Consider a communicatingclass $\mathcal{Y}_e$ without direct successor class, that is, a class at the end of the directed graph. Since no states outside of $\mathcal{Y}_e$ can ever be reached, the subfamily $\left(\mathcal{E}_y(r)\right)_{y\in\mathcal{Y}_e}$ of equilibrium payoff sets can be computed from the algorithm in Proposition 4.9 without considering states in $\mathcal{Y}_e^c$. This is particularly simple if $\mathcal{Y}_e = \{y_e\}$ is a singleton, that is, y_e is an absorbing state. The continuation game is then just a repeated game and hence $\mathcal{E}_{y_e}(r)$ can be computed with Theorem 2 in Sannikov [29]. One can then proceed backwards in the directed graph: consider a communicating class $\mathcal{Y}'$, for which all subfamilies of equilibrium payoff sets of direct successor classes $\mathcal{Y}_{e_1}, \ldots, \mathcal{Y}_{e_n}$ have been computed already. In the computation of $\left(\mathcal{E}_y(r)\right)_{y\in\mathcal{Y}'}$, incentives from state transitions to states in $\mathcal{Y}_E := \bigcup_{k=1}^n \mathcal{Y}_{e_k}$ do not need to computed iteratively as in Proposition 4.9, but only incentives via state transitions within the communicating class $\mathcal{Y}'$ have to be accounted for in an iterative fashion. This becomes again particularly simple if $\mathcal{Y}' = \{y'\}$ is a singleton. Then $\mathcal{E}_{y'}(r) = \mathcal{B}_{r,y'}\left((\mathcal{E}_y(r))_{y\in\mathcal{Y}_E}\right)$ can be solved in a single application of Theorem 6.7 rather than an iterated application. We refer to Section 8 in Bernard [4] for notes on the implementation of Theorem 6.7.

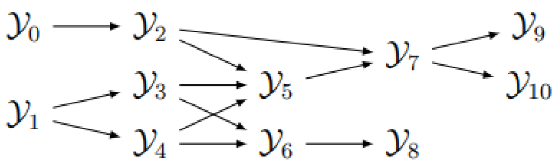

Figure 11: Communicating classes of a stochastic game form a directed acyclic graph.

How did the author get there? (2) Intuition

我覺得是有限的狀態，如此就可以用有限的步數計算完。