

MULTI-STAGE GAME THEORY IN CONTINUOUS TIME

BY IOANID ROSU¹

I define multi-stage stochastic games in continuous time. As in Bergin and MacLeod (1993), strategies have infinitesimal inertia, i.e., agents cannot change their strategies in an infinitesimal interval immediately after each time t . I extend the framework to allow for mixed strategies. As a novel feature in continuous time, mixing can be done both over actions, and over time (choosing the time of the action). I also define "layered times," which allow for stopping the clock and having various stages of the game be played at the same moment in time. I apply the theory to a trading game, where patient agents can choose whether to trade immediately or place a limit order and wait.

KEYWORDS: Continuous time game theory, mixed strategies, stopping the clock.

1. INTRODUCTION

The framework I use in this paper borrows mainly from the theory of repeated games in continuous time, as developed by Bergin and MacLeod (1993). I extend their framework by allowing stochastic moves by Nature, entry of new players, and mixed strategies. An alternative definition of continuous-time game theory can be found in Simon and Stinchcombe (1989). They define a continuous-time game to be a limit of discrete games, which makes their framework intuitive, but very hard to work with, especially when it comes to mixed strategies.

Date: January 2006.

¹The author thanks Drew Fudenberg, Sergei Izmalkov, Andrew Lo, Dimitri Vayanos, and Jiang Wang for helpful comments and suggestions. He is also grateful to participants at the NBER meeting, May 2004; WFA meeting, June 2005; and to seminar audiences at MIT, Carnegie Mellon, and U Chicago.

Continuous-time game theory is not a straightforward extension of discrete time game theory.² There are a few conceptual problems, as pointed out by Simon and Stinchcombe (1989), or Bergin and MacLeod (1993). To understand why, suppose one tries to replicate a typical punishment strategy from discrete time repeated games:

Continue to cooperate if the other player has not defected yet; if the other player defected at any point in the past, *immediately* defect and continue to defect forever.

The difficulty to make this strategy precise is two-fold. First, in continuous time there is no first time after t , which makes it difficult to “continue” a certain course. One way to get around this problem is to allow strategies to have inertia. But this creates a second problem, since the other players can take advantage of inertia.³ One way to allow players to react immediately is to enlarge the concept of strategy to include sequences of faster and faster responses. The mathematical concept that allows to do that is *completion* with respect to a metric (see below).⁴

Also, besides the usual problems with game theory in continuous time, there is an extra problem when dealing with multi-stage games. To wit, suppose an agent exits the game at time t . When is then the next stage of the game played? Since in continuous time there is no first time after t , one is compelled to have the next stage-game played also at t . I do this by introducing “layered times,” i.e., by allowing multiple games to be played at at the same time.

²For the discrete-time version of this theory, see Fudenberg and Tirole (1991), ch. 4.

³One could force the players to all have the same inertia, but then this would be equivalent to forcing the game to take place in discrete time.

⁴Another way to define immediacy is by using infinitesimal numbers, which is the mathematical field of non-standard analysis.

A third problem that arises in continuous time is that there is no first time before t . This issue is important when one needs a description of the game right before t . The solution of this problem is to allow only strategies that behave well immediately before any time t . The technical concept, inspired from Simon and Stinchcombe (1989), is of a strategy with a uniformly bounded number of jumps (to be defined below).

The last extension I consider in order to define multi-stage games in continuous time is that of mixed strategies. Unlike discrete time, in continuous time mixing can be done both over actions, and over time (choosing the time of an action).⁵

Finally, in this paper all information, together with agents' strategies and beliefs are common knowledge.

2. TRADING EXAMPLE

To better explain the intuition of this theory, I employ a trading game very similar to that in Rosu (2006). In this example, patient sellers endowed with one unit of the asset lose utility proportional to their expected waiting time. Sellers get a random liquidity shock, arrive to the market, and decide whether to sell immediately at the reservation price (and exit the market), or place a limit order and wait. Impatient buyers also arrive randomly at the market, and always place market orders against the sellers' limit orders.

Since this is a model of continuous trading, it is useful to set the game in continuous time. There are also technical reasons why that would be useful: in continuous time,

⁵It turns out that a strategy that mixes over time can be described as a limit of strategies that mix over actions — in the same way that a strategy with infinitesimal inertia is a limit of strategies with ε -inertia for ε a fixed positive number.

with Poisson arrivals the probability that two agents arrive at the same time is zero. This simplifies the analysis of the game.

Another important benefit of setting the game in continuous time is that agents can respond immediately. More precisely, one can use strategies that specify: “Keep the limit order at a_1 as long as the other agent stays at a_2 or below. If at some time t the other agent places an order above a_2 , then *immediately after* t undercut at a_2 .” Immediate punishment allows simple solutions, whereby existing traders do not need to change their strategy until the arrival of the next trader.

But there are problems that need to be fixed when setting the game in continuous time. Suppose a trader submits a market order at t and exits the game. The next stage of the game will then be played with fewer traders. But at which time will this next stage game be played? No $t + \varepsilon > t$ is satisfactory, because it would imply agents waiting for a positive time, during which they lose utility. The solution is to “stop the clock,” so that the next game is also played at t (hence the idea of “layered times”). The clock is restarted only when in the stage game no agent submits a market order.

Also, suppose that an impatient buyer arrives suddenly to the market and submits a market order at time t . What is the ask price at which that order is to be executed? Since there is no last time before t , one needs to have a well defined notion of the outcome of the game immediately before t . This means that one should use strategies that do not behave too wildly. The technical concept, inspired from Simon and Stinchcombe (1989), is of a strategy with a uniformly bounded number of jumps.

The last concept that is useful in this trading example is of a mixed strategy. It turns out that in equilibrium, in the state with the largest number of sellers, the seller with

the lowest offer has a mixed strategy. Mixing can be done over either actions or over time, but in this particular case mixing over time works better.⁶

3. MULTI-STAGE GAMES IN CONTINUOUS TIME WITH PERFECT INFORMATION

In order to define a game, one must define the spaces of actions, outcomes, and strategies. The definitions follow closely those of Bergin and MacLeod (1993). I extend their framework in several directions: (i) there is a well-defined description of the game right before any time t ; (ii) I allow for entry decisions of new agents; and (iii) I account for the possibility of having more than one game played at the same time. I start with an infinitely countable set of players I , but I assume that at each stage there are only finitely many agents in the game. (In the trading example this is true, since traders arrive according to independent Poisson processes, so with probability one at each point in time there are only finitely many traders.)

I want to include the case where at some times t the game is played more than once. I do this by taking the product of the time interval $[0, \infty)$ with the set of natural numbers \mathbb{N} , to indicate how many times a game has been played at some time t . Define

$$(1) \quad \mathcal{T} = [0, \infty) \times \mathbb{N}$$

the set of times at which players can move, counted with multiplicity. Notice that if \leq is the lexicographic order, (\mathcal{T}, \leq) is a totally ordered space. Denote the element $(0, 0) \in \mathcal{T}$ also by 0. Define intervals in \mathcal{T} in the usual way: for example, if $T = (t, n) \in \mathcal{T}$ define $[0, T) = \{T' \in \mathcal{T} \mid 0 \leq T' < T\}$. When there is no danger of confusion, write t instead

⁶There are equilibria with mixing over time for all values of the parameters, while equilibria with mixing over actions may not exist for some values of the parameters.

of (t, n) . Also, define a measure on $[0, \infty)$ so that bounded measurable functions are integrable:

$$(2) \quad \mu(dt) = e^{-at} dt.$$

In general, I want the action space for player i to be a compact complete metric space (X_i, d_i) . Typically, X_i is a compact subset of \mathbb{R}^n and d_i is the inherited metric. In the trading example, the action space for player $i \in I$ can be defined as a subset of \mathbb{R}^2 :

$$(3) \quad X_i = ([B, A] \times \{0, 1\}) \cup \{\mathbf{out}\},$$

where \mathbf{out} is some point in \mathbb{R}^2 which does not lie on $[B, A] \times \{0, 1\}$. An action $(x_i, 1) \in X_i$ is interpreted as a limit order at $x_i \in [B, A]$. An action $(x_i, 0) \in X_i$ is interpreted as a market sell (buy) order, in which case x_i is the current bid (ask) price, respectively. The action $x_i = \mathbf{out}$ indicates that either (i) player i has not entered the game yet; or (ii) player i exited the game before. One could also allow agents to exit freely at time t . This will not happen in equilibrium if the utility from exiting is very small, so in order to simplify the description of the game I do not allow free exit. Define also projections on the first and second factor, $\pi_1 : X_i \rightarrow [B, A] \cup \{\mathbf{out}\}$ and $\pi_2 : X_i \rightarrow \{0, 1\} \cup \{\mathbf{out}\}$, in the obvious way.

I now define outcomes of the game. Let \mathcal{B}_{X_i} and \mathcal{B}_X be the Borel sets of X_i and $X = \prod_{i \in I} X_i$, respectively; and let \mathcal{B} be the Borel sets of $[0, \infty)$. A function $\nu : [0, \infty) \rightarrow \mathbb{N}$ is said to have finite support if ν is zero everywhere except on a finite set $M_1 = \{t_1, \dots, t_K\}$ (its support). One also associates the set $M = \{(t_1, n_1), \dots, (t_K, n_K)\}$, where all $n_k =$

$\nu(t_k) > 0$. Vice versa, for any such set M one can define a function $\nu_M : [0, \infty) \rightarrow \mathbb{N}$ with finite support by sending $t \in [0, \infty)$ to zero if $t \notin M_1$; and to n_k if $t = t_k \in M_1$.

Definition 1. *Let X be a space with measure. A function $f : \mathcal{T} \rightarrow X$ is called layered if there exists a function $\nu : [0, \infty) \rightarrow \mathbb{N}$ with finite support such that $\forall t \in [0, \infty)$ and $\forall n, n' > \nu(t)$ one has $f(t, n) = f(t, n')$. If $f : \mathcal{T} \rightarrow X$ is layered, associate a function $f^\nu : [0, \infty) \rightarrow X$ by $f^\nu(t) = f(t, \nu(t))$. I say that f is a layered measurable function if f^ν is measurable. An outcome for player i is a layered Borel measurable function $h_i : \mathcal{T} \rightarrow X_i$.*

So an outcome is like a regular measurable function $h_i : [0, \infty) \rightarrow X_i$, except that at a finite set $\{t_1, \dots, t_K\}$ (the support of ν) it can take several values, up to the integer number $\nu(t_k)$. This corresponds to the idea that at some times t_k the game can be played more than once (in my case, if some agent places a market order).

I call the function ν the *layer* of f . Sometimes I also call the layer of f the associated set $M = \{(t_1, n_1), \dots, (t_K, n_K)\}$, with $n_k = \nu(t_k)$. Also, if f_1 and f_2 are two layered functions with layers ν_1 and ν_2 , one can take the combined layer of f_1 and f_2 to be $\nu = \max\{\nu_1, \nu_2\}$. This is useful for situations where one has to compare f_1 and f_2 . Consider a layer ν . Then I define: \mathcal{T}^ν , the set of layered times associated with ν ; H_i , the space of outcomes for player i ; and H_i^ν , the space of outcomes associated with ν :

$$(4) \quad \mathcal{T}^\nu = \{(t, n) \in \mathcal{T} \mid n \leq \nu(t)\},$$

$$(5) \quad H_i = \{h_i : \mathcal{T} \rightarrow X_i \mid h_i \text{ layered measurable}\},$$

$$(6) \quad H_i^\nu = \{h_i : \mathcal{T} \rightarrow X_i \mid h_i \text{ layered measurable with layer } \nu\}.$$

This is a metric space with the metric $D_i : H_i^\nu \times H_i^\nu \rightarrow \mathbb{R}_+$ given by

$$(7) \quad D_i(h_i, h'_i) = \int_{[0, \infty)} d_i(h_i^\nu(t), h_i^{\nu'}(t)) \mu(dt) + \sum_{k=1}^K \sum_{n=0}^{\nu(t_k)} d_i(h_i(t_k, n), h'_i(t_k, n)).$$

Rewrite this as

$$(8) \quad D_i(h_i, h'_i) = \int_{T^\nu} d_i(h_i(T), h'_i(T)) \mu^\nu(dT).$$

Since the space of measurable functions $f_i : [0, \infty) \rightarrow X_i$ is compact and complete, so is H_i^ν . Now, if $\nu \leq \nu'$, there is an inclusion $H_i^\nu \rightarrow H_i^{\nu'}$. Also, one knows that for every two layers ν_1 and ν_2 one can take their maximum $\nu = \max\{\nu_1, \nu_2\}$, which satisfies $\nu_1, \nu_2 \leq \nu$.

This means that one can regard H_i as the limit of H_i^ν when ν becomes larger and larger.

Because of this, H_i is a metric space, but it might not be either complete or compact.

I now define the space H of outcomes of the game. For this, let $H^\nu = \prod_{i \in I} H_i^\nu$ the product space with the metric $D = \prod_{i \in I} \frac{1}{2^i} D_i$. It is a standard exercise in measure theory to see that H^ν is compact and complete. As before, if $\nu \leq \nu'$, there is an inclusion $H^\nu \rightarrow H^{\nu'}$. I then define H as the union of all H^ν for all layers ν . This is still a metric space, but it might not be complete or compact. To justify this definition, consider an outcome $h \in H$. Since h belongs to a union of H^ν over all layers ν , there must exist a particular ν so that $h \in H^\nu$ (in which case, I say that ν is the layer of h). This corresponds to the fact that all agents are in the same game, played at the times described by the layer ν .

Also, if $Z \subset \mathcal{T}^\nu$ is layered measurable, and $h_i, h'_i \in H_i^\nu$, define a metric relative to Z by $D_i(h_i, h'_i, Z) = \int_Z d_i(h_i(T), h'_i(T)) \mu^\nu(dT)$. Define also a metric D on H relative to Z in the same way it was done for the product metric above.

Now I define strategies. In discrete time, pure strategies map histories to actions, while mixed strategies map histories to probability densities over actions. For technical reasons it is easier to think of a history as an outcome of the game together with a time t at which history is taken. This way, one can define a strategy as a map from $\{\text{outcomes} \times \text{times}\}$ to $\{\text{actions}\}$. Formally, a strategy for agent i is a map

$$(9) \quad s_i : H \times \mathcal{T} \rightarrow X_i$$

which satisfies the following axioms

A1. The function s_i is layered measurable on $H \times \mathcal{T}$.

A2. For all $h, h' \in H$ and $T \in \mathcal{T}$ such that $D(h, h', [0, T]) = 0$, one has $s_i(h, T) = s_i(h', T)$.

The second axiom ensures that future does not affect current decisions. Rewrite

$$h \sim_T h' \iff D(h, h', [0, T]) = 0.$$

As it was discussed above, these two axioms alone do not ensure that strategies uniquely determine outcomes. For that, one needs some inertia condition. If $t \in [0, \infty)$ and ν is a layer, denote by $t^\nu = (t, \nu(t))$, and $t = (t, 0)$.

A3. The function s_i displays inertia, i.e., for any $h \in H^\nu$ and any $t \in [0, \infty)$, there exists $\varepsilon > 0$ and $x_i \in X_i$ such that

$$D_i(s_i(h'), x_i, [t^\nu, t + \varepsilon)) = 0$$

for every $h' \in H^\nu$ such that $h \sim_{t^\nu} h'$.

Denote by S_i the set of functions s_i on $H \times \mathcal{T}$ which satisfy A1, A2, A3. Denote by $S = \prod_{i \in I} S_i$. The next theorem shows that a strategy profile $s = (s_i)_i$, i.e., a set of strategies s_i for for each player $i \in I$, uniquely determine an outcome on every subgame. More precisely one has the following result:

Proposition 1. *Let $s \in S$. Then for every $h \in H$ and $T \in \mathcal{T}$, there exists a unique (continuation) outcome $\tilde{h} \in H$ so that $h \sim_T \tilde{h}$ and $D(s(\tilde{h}), \tilde{h}, [T, \infty)) = 0$.*

Proof. The proof is the same as in Bergin and MacLeod (1993), but one has to make sure that one works in H^ν for some layer ν . □

Given $(h, T) \in H \times \mathcal{T}$ and $s \in S$, denote by $\sigma(s, h, t)$ the outcome which agrees with h on $[0, T)$ and is determined by the strategy s on $[T, \infty)$. Let $s_i, s'_i \in S_i$. I now define a metric on S_i :

$$(10) \quad \rho_i(s_i, s'_i) = \sup_{H \times \mathcal{T} \times S_{-i}} D(\sigma((s_i, s_{-i}), h, T), \sigma((s'_i, s_{-i}), h, T)).$$

One also has to introduce an axiom which ensures that for each t the outcome of the game right before t is well defined. One way of doing this is to restrict to strategies s_i that lead to locally constant outcomes with a uniformly bounded number of jumps.

A4. For the strategy s_i there exists M (depending only on s_i) such that for any strategies s_{-i} of the other players, the outcome $\sigma_i((s_i, s_{-i}), h, t)$ for player i is locally constant and has at most M jumps.

Redefine S_i to include on the strategies that satisfy A4. Now recall that at each $t \in [0, \infty)$ the strategies have inertia for some ε (depending on t). I want inertia to be infinitesimal, because I want to allow for immediate responses. This can be done by completing the space of strategies: Denote by S_i^* the completion of S_i with respect to the metric ρ_i , and by $S^* = \prod_{i \in I} S_i^*$. Completion is done so that the upper bound for the number of jumps is the same for all. More precisely, a point in S_i^* corresponds to a Cauchy sequence $(s_i^n)_n$ of strategies in S_i , and one demands that there exists M so that for each n , s_i^n jumps at most M times, regardless of the other players' strategies. The following result shows that to each strategy in S^* one can associate a unique outcome in H .

Proposition 2. *For every $s \in S^*$ and every $(h, T) \in H \times \mathcal{T}$, there exists a unique h^* such that $\sigma(s^n, h, T) \rightarrow h^*$ for any Cauchy sequence $(s^n)_n$ in S converging to s .*

Proof. If $(h, T) \in H \times \mathcal{T}$, there exists a layer ν so that $h \in H^\nu$ and $T \in \mathcal{T}^\nu$. The result then follows easily since H^ν is compact and complete. \square

I have just showed that for $s \in S^*$ one can associate a unique outcome of the (whole) game, which I denote by $\sigma^*(s)$. Because completion is done using the same upper bound for the number of jumps, the following result is straightforward. The result allows one to talk about the outcome of a game right before some time t .

Proposition 3. *The outcome $\sigma^*(s)$ associated to a strategy $s \in S^*$ is left-continuous.*

I am almost done in defining the game. The only thing that is left is to describe the payoff for some strategy $s \in S^*$ in a subgame defined by a history $(h, T) \in H \times \mathcal{T}$. Since strategies uniquely define outcomes in every subgame, as long as there exists some payoff $u_i(\sigma^*(s, h, T))$ for each agent i . Define now the equilibrium concept:

Definition 2. *A strategy profile $s \in S^*$ is an ε -Nash equilibrium (ε -NE) if for any $h \in H$*

$$(11) \quad u_i(\sigma^*(s, h, 0)) \geq u_i(\sigma^*((s'_i, s_{-i}), h, 0)) - \varepsilon, \quad \forall i \in I, \quad \forall x'_i \in S_i^*.$$

A strategy profile $s \in S^$ is an ε -subgame perfect Nash equilibrium (ε -SPE) if for any $(h, T) \in H \times \mathcal{T}$*

$$(12) \quad u_i(\sigma^*(s, h, T)) \geq u_i(\sigma^*((s'_i, s_{-i}), h, T)), \quad \forall i \in I, \quad \forall x'_i \in S_i^*.$$

For $\varepsilon = 0$ in the above inequalities one obtains the concepts of Nash equilibrium (NE) and subgame perfect Nash equilibrium (SPE). One has the following important result.

Proposition 4. *A strategy profile $s \in S^*$ is a subgame perfect equilibrium if and only if for any Cauchy sequence $(s^n)_n$ converging to s , there is a sequence $\varepsilon^n \rightarrow 0$ such that s^n is an ε^n -subgame perfect equilibrium.*

Proof. The proof is essentially the same as in Bergin and MacLeod (1993), but again one has to make sure that one works in H^ν for some layer ν . □

I now discuss mixed strategies. For simplicity of discussion I omit the presence of layers, so one takes $\mathcal{T} = [0, \infty)$. Consistent with this philosophy of locally constant

outcomes and inertia strategies, I want to have mixed strategies randomly switch over a small interval. More formally, let X_i be the space of actions for player i , and $[0, \infty]$ the metric space with metric $d(x, y) = |e^{-x} - e^{-y}|$. Define a mixed strategy to be a measurable function

$$(13) \quad s_i : H \times \mathcal{T} \rightarrow X_i \times X_i \times [0, \infty],$$

where the first component of s_i is the initial action in X_i ; the second component is the action to which s_i will switch in the interval of time right after t ; the third component is the Poisson intensity of switching. I call this type of strategy “mixed over time,” because randomness only comes from the time of switching, while the actions before or after switching are deterministically chosen. One can also allow for mixing over actions, in which case one has to replace X_i by $\Phi(X_i)$, the set of probability densities over X_i , i.e., the set of non-negative integrable functions on X_i with total integral equal to one. Since $\Phi(X_i)$ is compact if X_i is compact, the analysis is essentially the same.

I say that the strategy s_i has inertia in a similar way as before, but one adds the requirement that, after switching, the action to which player i switched will be also held constant for a small period of time. Then one has to modify the description of outcomes, which are now stochastic processes that are built in a very similar way to Poisson processes. The space of strategies is also constructed by taking a completion, in the same way it was done for pure strategies.

To see how a strategy mixed over time works, consider the general case of the trading example, where Nature moves at each time t by bringing new players to the game. One can consider in fact two instances of Nature (two distinct players), each of whom

brings one type of trader to the game. The space of actions for each instance of Nature, e.g., the one that brings patient sellers, is the set 2^I of all subsets of I (in principle I allow Nature to add or remove any players from the game). Nature plays the following strategy: if $(h, t) \in H \times \mathcal{T}$ is the history at t , and $J = I_{t-}$ is the set of players in the game right before t , then $s_N(h, t) = (J, J \cup \{PS\}, \lambda_{PS})$.

I now briefly discuss the notions of equilibrium that can be defined in this framework. The notions of subgame perfect equilibrium and Markov (perfect) equilibrium are simple extensions of the corresponding concepts in discrete time (see Fudenberg and Tirole (1991)). One can also define the notion of *competitive* Markov equilibrium. This is a Markov perfect equilibrium from which local deviations can be stopped by local punishments, assuming behavior in the rest of the game does not change. In other words, if one truncates the equilibrium strategies by looking at some time interval $(t, t + \delta)$ (truncation is possible because of the Markov condition), the restrictions of the strategies to this time interval remain a Markov equilibrium.

REFERENCES

- [1] BERGIN, J., AND B. MACLEOD (1993): “Continuous Time Repeated Games,” *International Economic Review*, 34, 21–37.
- [2] BIAIS, B., D. MARTIMORT, AND J-C. ROCHET (2000): “Competing Mechanisms in a Common Value Environment,” *Econometrica*, 68, 799–837.
- [3] FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*, MIT Press.
- [4] ROSU, I. (2006): “A Dynamic Model of the Limit Order Book,” Working Paper, University of Chicago.
- [5] SIMON, L., AND M. STINCHCOMBE (1989): “Extensive Form Games in Continuous Time: Pure Strategies,” *Econometrica*, 57, 1171–1214.

A game-theory estimate is made of the effectiveness of the predistortion and correction of signals on the basis of the minimax criterion of the average risk. The cost of the game (minimax average risk) is determined, together with the optimal strategies for operators in systems for transmitting information and radio-interference systems. Read more. Article. In this paper, credibility is studied for two-person, single-stage games. A strategy is said to be fully credible if it satisfies the Principles of [Show full abstract] Optimality under the Stackelberg setting. Multi-Stage Games. This is page 1 Printer: Opaq. Recall that the normal form game had players choose their actions simultaneously, and ignored any real issues concerning sequential behavior. In reality, dynamic play over time may be more complex than one game that unfolds over time. Instead, players can play one game that is followed by another, or maybe even several other games.