4109H: SINGLE DEVIATION PRINCIPLE

Sasha Vostroknutov

In this short note I try to explain how single deviation principle works. I have to use some mathematical notation, which I did not use in class, in order to be precise about things.

Let $G = \{A_1, A_2; u_1, u_2\}$ be a normal form game with 2 players who have action sets A_1 and A_2 and payoffs u_1 and u_2 (remember, u_i is a function from $A_1 \times A_2$ to $\mathbb{R}$). Let $G(T)$ be finitely repeated game: G is repeated T times, the payoffs in the end are the sum of payoffs in each stage. Let $G(\delta, \infty)$ be infinitely repeated game: the payoffs are the δ-discounted sum of payoffs in each period.

Denote by $o_t = (a_{1t}, a_{2t}) \in A_1 \times A_2$ the outcome of playing G in period t. Here $a_{it} \in A_i$ for $i = 1, 2$. The history of play after t periods is then

$$h_t = (o_1, o_2, ..., o_t)$$

for some o_k, $k = 1..t$. Notice that h_t contains the choices of *both* players in all t periods.

Denote by H_T the set of all possible histories of lengths from 1 to $T - 1$ and by H the set of all possible histories of any length.

What is the strategy of player i in $G(T)$? Player i has to specify what he is going to do after each possible history. Therefore, the strategy is a function $s_i : H_T \to A_i$. For each element $h \in H_T$, which is the history of some length, s_i specifies the choice of action $s_i(h) \in A_i$ that will be played in the period following h. In the same manner, the strategy of player i in $G(\delta, \infty)$ is the specification of the action after each possible history of any length. So, it is a function $f_i : H \to A_i$.

Now, given any s_i, I want to define what is a *single deviation from s_i at history $\hat{h}$*. In words, single deviation from s_i at $\hat{h}$ is any other strategy $\hat{s}_i : H_T \to A_i$ which is different from s_i in exactly one place, namely after history $\hat{h}$. This means that $\hat{s}_i(\hat{h}) \neq s_i(\hat{h})$ and $\hat{s}_i(h) = s_i(h)$ for all other $h \in H_T$.

To deliberate more on the difference between s_i and $\hat{s}_i$ let's consider the following exercise. We decide to compare what player i is doing if he follows s_i or $\hat{s}_i$ after any possible history. If we pick any history $h \neq \hat{h}$ and see what the two strategies prescribe, we will not be able to detect the difference. Only if we pick $\hat{h}$ we will be able to see the difference in choices. Notice that even if we pick a history h that follows $\hat{h}$ *we will still not be able to see the difference*. In yet other words, if you think of strategies s_i and $\hat{s}_i$ as choices in all information sets of player i on the tree of $G(T)$, then the strategies will be different in exactly one information set.

The same definition can be given for the single deviation from f_i in $G(\delta, \infty)$: In the two paragraphs above replace H_T with H, s_i with f_i and $\hat{s}_i$ with $\hat{f}_i$.

Now we are ready to state the single deviation principle:

Proposition 0.1 *The strategies s_1 and s_2 of players 1 and 2 playing $G(T)$ constitute SPNE if and only if for any history $\hat{h} \in H_T$ and any $i = 1, 2$ player i has no profitable single deviation from s_i at $\hat{h}$ given that $\hat{h}$ is reached and given that s_{-i} is fixed.*

The same works for infinitely repeated game:

Proposition 0.2 *The strategies f_1 and f_2 of players 1 and 2 playing $G(\delta, \infty)$ constitute SPNE if and only if for any history $\hat{h} \in H$ and any $i = 1, 2$ player i has no profitable single deviation from f_i at $\hat{h}$ given that $\hat{h}$ is reached and given that f_{-i} is fixed.*

Pay careful attention to the last phrase of these propositions. We require that there is no single deviation after any history *given* that it is reached. This means that we have to check all possible histories even if they can never occur under the equilibrium strategies. This is in accord with SPNE which requires that nobody wants to deviate in any subgame regardless if it is reached or not.

In addition, notice that these propositions are of “if and only if” type. What does this mean? First, if something is SPNE then there are no single deviations anywhere (which is not that particularly interesting). Second, if there are no single deviations then the strategies are SPNE. Now, this is interesting because this can be used to find SPNE.

I understand that it is probably tough for you to understand all these abstract definitions. Unfortunately I could not find any other way to express these ideas precisely. However, single deviation stuff can simplify your life considerably when it gets to proving whether something is SPNE or not. Think about it: without the principle in order to show SPNE you need to consider *all possible deviations after all possible histories.* This is no trivial task since deviations might happen not only in one place but in say 6, 10, 1000 or even in infinitely many places. If you are dealing with strategies a bit more complicated than trigger, you existence can become extremely miserable.

The way to get used to the single deviation principle is to consider some examples, which turn out to be not that complicated.

Example. Tit-for-Tat against itself in infinitely repeated Prisoner's Dilemma. Consider the PD game:

	C	D
C	4, 4	0, 5
D	5, 0	1, 1

TfT strategy says:

Play C in the first period. In period $t > 1$ play whatever the other player did in period $t - 1$.

Notice that this is a complete description of the strategy, which should tell us what the player does after each possible history.

Let us show that in an infinitely repeated PD it is not SPNE for both players to play TfT if δ is high enough. Look at the definitions of the single deviation principle. It says that to show that something is SPNE we need to show the absence of single deviations for all histories. But then, to show that something is *not* SPNE it is enough to find *one* history at which for some player there is a deviation! Let's see what this history can be.

First notice that on the equilibrium path both players always play C, repeating previous action of the other player. Now, look at how two TfT players behave after the history that ends with the outcome (C, D). If they both stick to TfT we will have the path

$$...(C, D) \mid (D, C), (C, D), (D, C), (C, D), (D, C)....$$

The continuation payoff for the player 1 here is

$$5 + 0 \cdot \delta + 5\delta^2 + 0 \cdot \delta^3 + 5\delta^4 + ... = \frac{5}{1 - \delta^2}$$

But what if player 1 does single deviation after the history ending with (C, D)? This means that instead of repeating D of player 2 he plays C and then continues doing whatever TfT was prescribing. We fix player 2's strategy and get the following path:

$$...(C, D) \mid (C, C), (C, C), (C, C), (C, C), (C, C)....$$

This gives player 1 the payoff

$$4 + 4\delta + 4\delta^2 + 4\delta^3 + ... = \frac{4}{1 - \delta}$$

Let's see for which δ it is profitable to deviate:

$$\begin{aligned} \frac{4}{1-\delta} &> \frac{5}{1-\delta^2} \quad \Rightarrow \\ \delta &> 1/4 \end{aligned}$$

For any δ bigger than $\frac{1}{4}$ the guy will prefer to deviate. So, TfT against TfT is not SPNE for high δ.

My advice to you: to prepare for the midterm, try to play with some 2x2 games and some simple strategies. Show that something is SPNE by considering all possible histories and showing that there are no profitable single deviations for high enough δ. Show that something is not SPNE by finding one history at which there is profitable deviation for high enough δ.