The following introduction serves as a simple explanation of the relevant theoretical concepts to be used to interpret the Honami & Kiyotaka Alliance Strategy. *Some of the "concepts" mentioned may look like a tangent at first.*

Roughly speaking, Game theory models various strategic interactions ("games") by splitting them into two different groups: one-shot (played once) and iterated (the same stage game played multiple times). One-shot versions have no prior history to condition on. However, strategies used in repeated games can take into account past plays.

That's not only categorization, but it's also one of the most important distinctions. However, it is more important that both versions have equilibrium with inherently incompatible properties.

Here is the typical payoff matrix for Prisoner Dilemma (PD), the almost canonical example used in Game theory.

Honami's class/Kiyotaka's class	Kiyotaka *cooperates*	Kiyotaka *defects*
Honami *cooperates*	R_H, R_k	S_H, T_k
Honami *defects*	T_H, S_k	P_H, P_k

P (both defects) > R (both cooperate) and T > S (one defect, the other cooperates) imply that the dominant strategy is to defect (for both or one player). But, as shown in the example below, it's usually more profitable when the opponent cooperates while you defect (strictly speaking, it's the mandatory condition for PD (T > R); however, since the interpretation shouldn't solely rely on a specific game, it makes sense to use a more vague condition).

Honami's class/Kiyotaka's class	Kiyotaka *cooperates*	Kiyotaka *defects*
Honami *cooperates*	3, 3	-1, 4
Honami *defects*	4, -1	1, 1

The rare example when cooperation is more profitable for both players in COTE is the Y2V10 exam for Honami and Kakeru. The main difficulty they faced was determining whether the opponent is trustworthy.

However, usually, for one-shot social dilemmas as well as PD, defection is *individually* optimal. It mainly happens because a game result has no consequence beyond the game context. There is no way to reward or punish the other player (using “reward” and “punishment” in the broadest sense). Simply put, no future = no leverage.

However, in the iterated versions, including the Iterated Prisoner Dilemma (IPD), the dilemma repeats over multiple rounds, allowing players to adjust strategies based on past interactions. This mirrors relationships, where trust and betrayal are not one-time events but ongoing dynamics.

Political scientist Robert Axelrod set out to identify strategies that perform well in IPD. He held a computer tournament to explore strategies for IPD (in which submitted programs played round-robin IPD matches). The result was striking. It turned out that the success lies in *cooperation*. Axelrod’s tournaments revealed that being *nice*, *forgiving*, *retaliatory*, and *not too clever* are fundamental for fostering cooperation. Despite attempts at clever manipulation, simple strategies like *Tit for Tat* (TFT) consistently triumphed, proving that in the game of trust, straightforwardness pays off.

The conditions necessary for a strategy to succeed in iterated games:

- **Nice.** Don’t defect first.
- **Retaliating.** Respond to defections with defection; otherwise, you’ll be exploited.
- **Forgiving.** Resume cooperation after defection or if the opponent does.
- **Don't be too clever.** The behavior should be easily understood to not cause retaliations on “false defections.”

However, Axelrod’s results are “empirical” and operate under several important restrictions. They are sound, yet they don’t fully picture the real-life cases.

For example, TFT is N*ash equilibrium* (NE), which, in simple words, is when nobody can do better or there is no initiative to deviate from the selected strategy, only under some conditions: infinitely repeated games, perfect monitoring, and a high *discount factor* (how much future payoff is valued relative to the current one), i.e., patient opponents. TFT is *generally* not subgame-perfect.

Actually, the problem with “infinitely repeated games” is the “end game” problem. Consider the situation when you know how many rounds you’re going to play. As soon as the last round is reached, the defection is more profitable; hence, there is no reason to cooperate. It's applicable for all players.

However, using simple backward induction, the same could be said about the round before last and so on up to the first round. Therefore, in practice, the last round should remain unknown. This condition is nearly sufficient. With these restrictions, cooperative strategies form NE.

Regarding *evolutionary stable strategy* (ESS), defined as a strategy that, once accepted by the majority in a population, cannot be invaded by a small group of mutants using a different strategy, the situation becomes even more complex. There is no pure strategy, which is ESS. Under some parameters, Generous TFT or Pavlov outperforms TFT (both are cooperative). However, generous cooperative strategies tend to survive and dominate. Though nasty strategies may occasionally appear, they do not survive and cannot dominate.

Another problem is the noise (anything that might affect links between intended and observed actions and decisions, such as implementation/observation errors). The most influential and widely known critique belongs to *Taleb*.

He presents it in his books for a broad readership and in his scholarly studies. His critique can be summarized as: don't over-rely on point equilibria or Gaussian (thin-tailed) distributions in thick-tailed or multimodal environments. His decision is to mitigate fragility by not relying on theatrical foresight ("they can see 9 moves ahead") and improve robustness through iterative adjustments driven by feedback with short steps.

Roughly speaking, the game theorists' reply focused on incorporating these points and developing methods to mitigate the risks. For instance, developing methods that don't rely on known distribution or build equilibria based on the worst-case outcome. The point is that Equilibria-like approaches remained relevant.

Folk theorem

There are multiple variations of this theorem (folk theorems are a class of theorems, actually), which were published for different games. However, the "intuitive" version of the theorem should be enough.

Intuitive version: credible reward-punishment schemes let participants use long-run losses as threats that outweigh any short-run gain from deviating.

Zero-determinant (ZD) strategies

In 2012, Press & Dyson published their article in which they introduced a special type of strategy with memory that allowed extortion, or, more broadly, allowed players to unilaterally impose linear relationships onto their opponents.

However, further research found out that extortion isn't ESS. Furthermore, it works only against naive opponents in the long run in noisy environments. *Generous* variants still can be stable and much more robust. For example, Stewart & Plotkin (2013) asserted, "*...generous ZD strategies are not only robust to being replaced by other strategies but can selectively replace any noncooperative ZD strategy. Generous strategies can be generalized beyond the space of ZD strategies, and they remain robust to invasion.*"

Law of requisite variety (Ashby's law)

In cybernetics, the term "variety" means the total number of distinguishable elements, for example, "inputs" used by a regulator. The law of variate states that a deterministic strategy can limit variety in outcomes by a finite value. Only adding variety to strategy can diminish variety in outcomes.

As a simple slogan, Ashby's law can be worded as "*only variety can destroy variety*."