Chapter Four: Incomplete Information Games*

4.1 The Concepts of Incomplete Information

Information is almost always incomplete in real life. But this does not mean that the issue of incomplete information should always be central to game theoretic modeling. It all depends on what question one wishes to address. Incomplete information is often about the players' true preferences, although it can also be about the players' true capabilities. In the first case, the uncertainty is about payoffs and whether their variation can affect the responses of the various "types" of players to the unfolding of the game. If that is the case, incomplete information models are best. But if the payoff variations are so limited that they cannot motivate drastic changes in choices between types, then complete information models are usually sufficient to address the main strategic issues.

The modern approach to incomplete information in Game Theory is due to John Harsanyi. One simply defines a one-shot game and duplicates it to account for the various types of players. One then adds information sets between matching turns for the player who lacks the relevant information. *Game 4* illustrated in Figure 1.8 (in Chapter 1) is perhaps the simplest example: Red is the player lacking information about Blue's type. And the consequences need not be just about Blue's payoffs since the outcome of facing the "wrong" type may also affect Red's payoffs. The most important aspect of Harsanyi's approach is in the treatment of beliefs at the information sets representing the uncertainty. One could simply set arbitrary probabilities on the various possible types. But that would ignore the fact that choices may inform on the type of the player who makes them. So, Harsanyi proposed to set arbitrary *initial* chances of facing one type or another and to update these chances according to what choices the corresponding types would make. In essence, if a type is more likely to make one choice rather than another, the very observation of that choice makes it more likely that one faces that type.

The idea can be formalized as follows: if there are n types $j \in \mathcal{J} = \{1, 2, .., n\}$ of player Blue denote by p_j the initial probabilities of facing a given type $j \in \mathcal{J}$.[1] And if $\mathbb{P}(\alpha|i)$ is the strategy-induced probability of the choice α being made by type $i \in \mathcal{J}$, then the "updated" probability b_i that Red is facing type i once α has been observed is

$$b_i = \mathbb{P}(i|\alpha) = \frac{\mathbb{P}(\alpha|i)p_i}{\sum\limits_{j=1}^{n}\mathbb{P}(\alpha|j)p_j} \qquad (1)$$

Formula (1) is called "Bayes' Law" and probability b_i is the "updated belief" about type i after choice α has been observed. Of course, beliefs can be updated again following another revealing move β. But the important point is that the other player(s), here Red, will now calculate her expected payoffs according to updated rather than static beliefs about Blue.

Of course, there are instance where a move α cannot possibly be optimal for any of the types with positive beliefs. In that case, that move and its possible consequences are "off-path"

[1] So $p_i \geq 0$ and $\sum\limits_{i=1}^{n} p_i = 1$.

according to the corresponding strategies and beliefs cannot be updated by Bayes' Law. It is then standard practice to set them arbitrarily.[2]

The central solution concept in games of incomplete information is the Perfect Bayesian Equilibrium (PBE): a PBE is a pairing of a strategy profile and a set of beliefs such that strategies are sequentially rational given the beliefs and the beliefs are updated according to Bayes' Law given strategies whenever possible, and are arbitrary for off-path information sets. There always exists such an equilibrium in any finite game of incomplete information. Figure 1.18 in Chapter 1 shows one PBE for Game 4 of Figure 1.8 and Figure 1.26 shows another one. The first PBE displays Bayesian updating of beliefs while the second shows the case of arbitrary beliefs for an off-path information set.

4.2 Pooling and Separating Equilibria

Equilibria can be of three different kinds at each information set. They can be pooling, separating or semi-separating. Figures 4.1, 4.2 and 4.3 are solutions of Game 4, from Chapter 1, with slightly changed payoffs and initial beliefs.

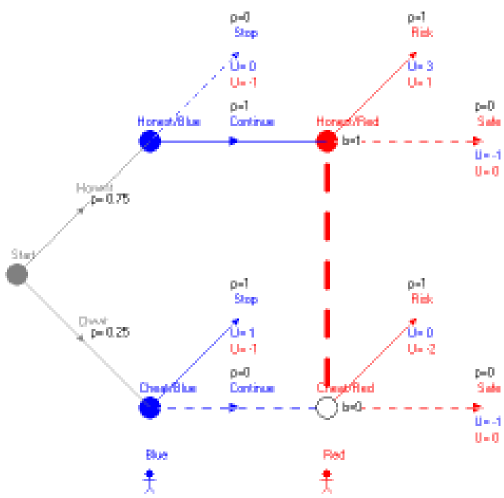

Figure 4.1: A Separating Equilibrium

Figure 4.1 illustrates a separating equilibrium: initial beliefs and payoffs are such that optimal play will reveal Player Blue's type to Red at her turn of play.[3] Figure 4.2 shows a pooling equilibrium of the same game that results from exchanging the Cheat Blue type's payoffs between Stop and Risk. In this case, Red cannot infer anything about Blue's type from his optimal play.

Finally, Figure 4.3 shows a semi-separating equilibrium. Here, the Cheat Blue type only chooses Continue with probability $p = \frac{1}{6} \simeq 0.16667$. Red cannot completely infer Blue's type from optimal play but can at least update her beliefs to much higher chance $b = \frac{2}{3} \simeq 0.66667$ that she is facing the Honest Blue type. The separation of types is not complete but substantial enough that it yields a high chance of not getting cheated. Since payoffs are the same in these last

[2]An interesting alternative is to obtain off-path beliefs by a "trembling" process that allows any move to have a tiny probability that is made to approach zero in a limit process. The resulting solution concept is called a "Sequential Equilibrium."

[3]There is another PBE where Blue always chooses Stop and therefore Red never gets a turn.

two figures, it is really the opposite initial beliefs that yield the different behavior, not just for Blue, but also for Red who can choose Risk with little worry in the first case, but only with probability $p = 0.5$ in the latter case.

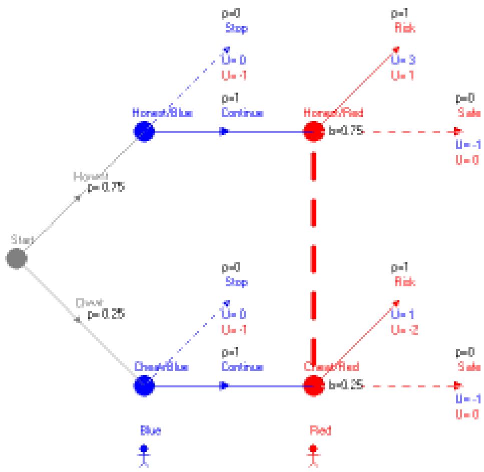

Figure 4.2: A Pooling Equilibrium

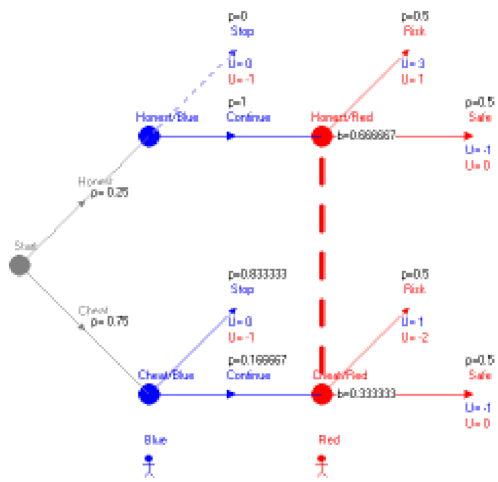

Figure 4.3: A Semi-Separating Equilibrium

The above features provide the technical means to develop strategies as well as mechanism designs to screen out undesirable types or to signal one's type in order to reach desirable outcomes. The next section discusses two typical examples.

4.3 Signaling and Screening

4.3.1 Getting a Promotion

Bill would like to get a promotion and is planning to go ask his boss Rachel. He has been looking around the job market for an alternative and has some reasonable leads. The question for Bill is how to approach Rachel: should he boast of some strong leads or simply state that he has been talking to other employers? Boasting may impress Rachel who obviously would not like to lose Bill, a performing employee, to a rival firm. But it may backfire and anger Rachel if she believes that Bill is bluffing about the strength of his leads. However, not boasting might leave Rachel plenty of room to turn him down, perhaps arguing that he needs a bit more experience in his current rank before aspiring for a higher one. It the end, it may all depend on how much Rachel believes his chances are out there. Bill, who has training in Game Theory, designed the following model:

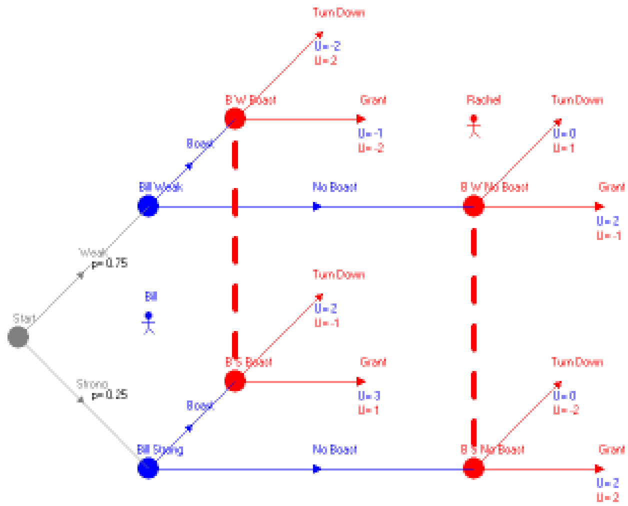

Figure 4.4: Promotion as a Signaling Game

Bill gives it only a 25% chance that Rachel believes he has a strong lead. He wonders how boasting would affect her decision. GamePlan indicates that this might be the right strategy if he does happen to have a strong lead. Indeed, there exists a single BPEs illustrated in Figure 4.5:

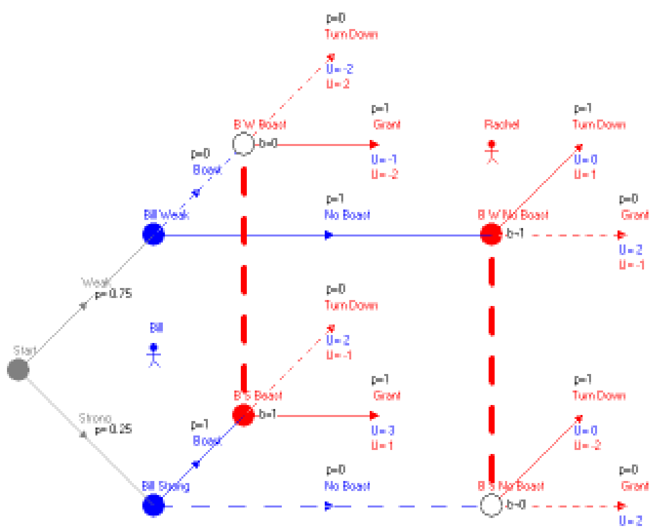

Figure 4.5: Signaling with a Separating PBE

But this outcome hinges on the payoff Bill assigned himself to obtaining a promotion under false pretenses: he granted himself $U = -1$ from the sequence {Weak, Boast, Grant} and $U = 0$ from the sequence {Weak, No Boast, TurnDown}. In other words, he couldn't live with lying. Shouls he be a litlle less honest with himself, the situation would change a bit. If one just exchanges the above two payoffs, one gets the result in Figure 4.6. Now, Bill can contemplate boasting with a weak hand but with a quite small probability $p = \frac{1}{6}$.

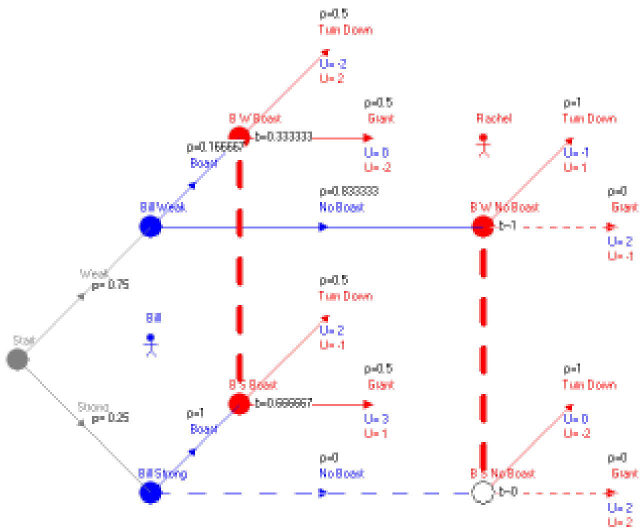

Figure 4.6: Signaling with a Semi-Separating PBE

The situation changes again if Rachel has different initial beliefs. If she thinks highly of Bill's chances by placing an initial probability of 75% that he has strong leads, then the solution shown in Figure 4.7 arises:

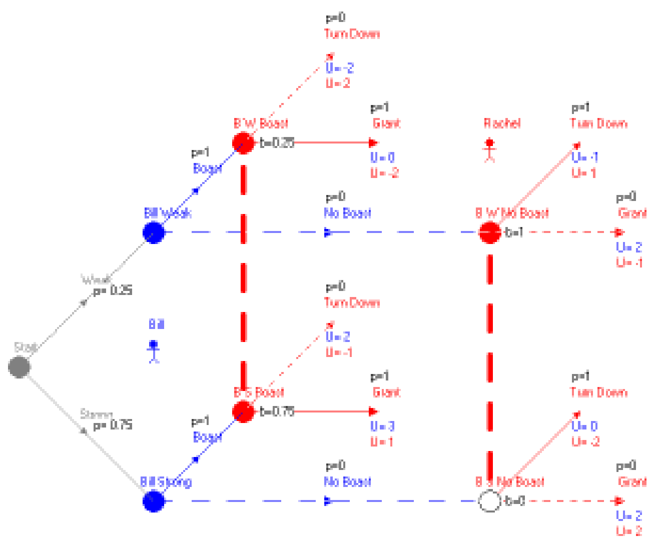

Figure 4.7: Taking Advantage of a Pooling Equilibrium

Now, Bill is always best off boasting and will always get his promotion, regardless of the strength of his leads. One can interpret that change as an effect of reputation. If Bill is perceived as a strong candidate he should act accordingly and will get his way.

4.3.2 Customer Service

One of the most frustrating issues of good customer service is the possibility of facing unreasonable or abusive customers. How does one distinguish between a honest customer complaining about a lemon and a dishonest one who simply decided s/he did not like the color of the purchased item after all? This will be discussed in the case study "Customer Service at B&S."

4.4 Deterrence

Deterrence is the art of persuading one's opponent to refrain from engaging in behavior one finds undesirable or harmful. There are two main avenues to achieve deterrence: (1) by threatening specific and credible retaliation; and (2) by establishing a reputation for toughness. The first avenue is studied in the context of repeated games in Chapter 5. The second can be discussed within the context of incomplete information.

4.4.1 Nuclear Deterrence

Hermann Kahn's description of a nuclear crisis as a Game of Chicken was briefly discussed in Chapter 3. It's representation as a normal form game is at best perfunctory. Chicken is really a game of timing and of deliberate escalations by each side willing to move closer to the abyss, defying the other side to take the next step toward mutual doom. Robert Powell was first to explore this idea in terms of an extensive form game. The basic building block is a sequence of alternative turns where each side can either quit or escalate, meaning getting closer to mutual doom. Doom is not a sure thing. Instead, there is an increasing probability that Nature will push the two sides over the top into a mutually disastrous outcome. An example is pictured in Figure 4.8. The probability of Doom increases by 0.1 with each escalation until a final Escalate would make it a sure thing. The last Escalate move shown in Figure 4.8 should in fact lead a further chance node and further turns for both sides until escalation yields Doom with certainty.

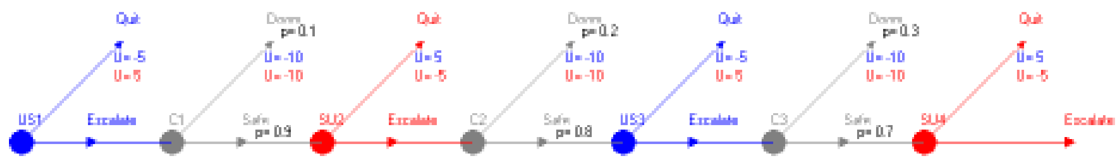

Figure 4.8: A Complete Information Chicken

Solving that game uses a typical backward induction. Depending on the expected payoff of reaching Doom, one side or the other will quit and the next to last will prevail. The game becomes interesting when the payoff of prevailing or quitting is uncertain for at least one side. Then, one Escalate move will certainly yield an expected payoff of Doom that exceeds the cost of quitting, but it is not clear who will encounter that calamity first. The game pictured in Figure 4.9 is an example. US is uncertain about SU's value for prevailing. If SU is Strong it values winning of losing (6 or -6) more highly than if it is Weak (5 or -5.) The values of the final Escalate move are derived from assumptions about US's willingness to escalate further. With 50% initial beliefs the game has a single PBE pictured below.

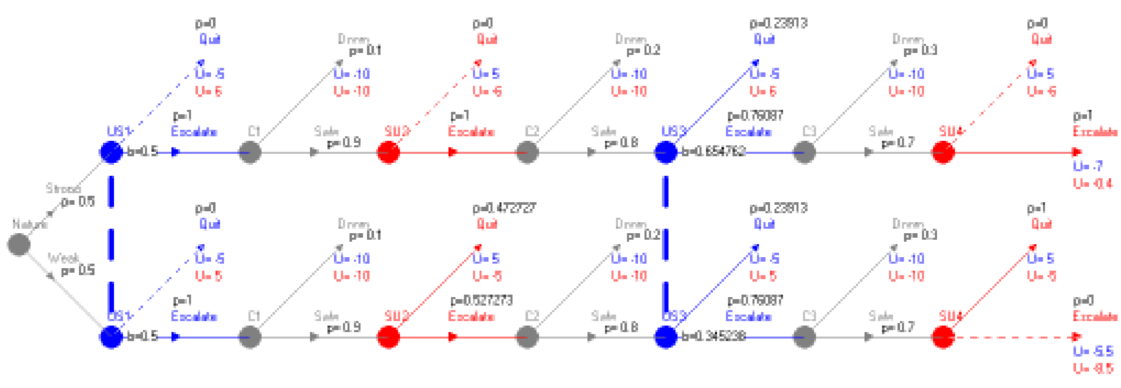

Figure 4.9: An Incomplete Information Game of Deterrence

US initially challenges SU which responds by a sure escalation when it is strong and a probabilistic one when it is weak. Facing escalation, US adjusts its beliefs upward on the chance of facing a strong SU and responds accordingly with a probability of escalating further. If the initial belief of facing a strong SU is increased enough, the US finds itself deterred from escalating in the first place. This is deterrence by reputation.

4.4.2 Entry Deterrence

Sometimes, reputation is established by prior behavior in similar circumstances. An example is a business facing the possibility of the multiple entry of competitors in its market. If it fights the first entrant despite high costs of doing so, it might even deter a potential second entrant from trying. An entry detrrence game is pictured with two typical solutions in Figures 4.10 and 4.11.

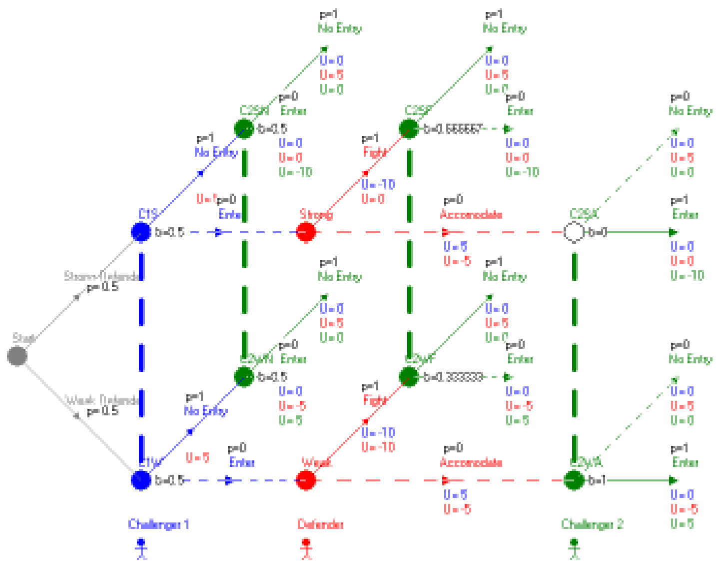

Figure 4.10: A Successful Deterrence of Entry

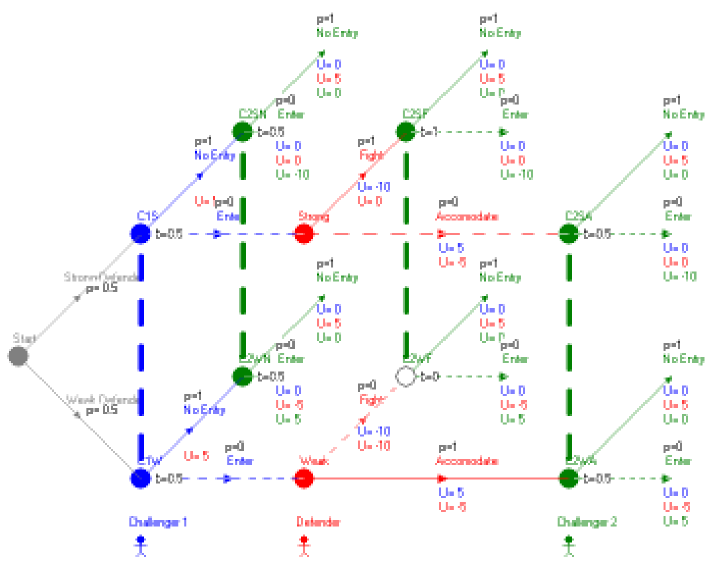

Figure 4.11: Another Successful Deterrence of Entry

The Defender in red faces two successive possible entrants: Challenger 1 who moves first followed by Challenger 2. The Defender can either Accomodate the first entrant and face a possible second competitor. Or it may Fight the first challenger and thereby send a signal of strength to the second challenger. In this model, the end moves at the green nodes have similar payoffs but are modified by intermediate payoffs on non-final moves, a useful GamePlan feature. With the given initial beliefs, one obtains several PBEs. The two solutions pictured illustrate successful deterrence. It is noteworthy that the second one has a signaling feature: the Defender would Fight if Strong but Accomodate if Weak.

4.5 Multiplayer Uncertainty

Many instances of incomplete information involve more than two players and beliefs strongly affect the outcome. Here again, we proceed by examples.

4.5.1 Alliance

History is full of examples where a challenger faced a defender supported by an ally. Arguably, the NATO alliance provided such support during the Cold War to West Germany and to Turkey facing threats from the Soviet Union. These were successful cases. A less successful case is examplified by the July Crisis leading to WWI: After the assassination of the heir to the Austrian throne in Sarajevo, Austria issued an unacceptable ultimatum to Serbia. Serbia turned to Russia, its ally, for protection. Russia turned to France, its ally, for support. Austria turned to Germany, its ally, for reassurance. And everybody turned to Britain for re-insurance... Germany issued a full private commitment to Austria to come to its defense should war break out. Emboldened by this strong support, Austria faced Russia down and things escalated out of control. Figures 4.12 and 4.13 show the solutions of a simple example of this kind of situation. The ally (in green) may or may not give a private commitment to the defender (in red.) The challenger (in Blue) makes decisions with that uncertainty.

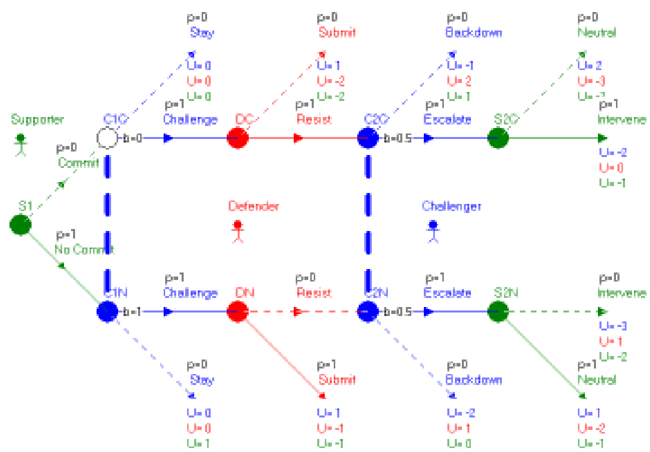

Figure 4.12: A Failure of Alliance Deterrence

In the first case, the supporter fails to issue a commitment and the challenger moves to test the defender's information. Since the challenger did not commit the defender submits immediately. Note that the information set {C2C, C2N} is off path and the beliefs there are arbitrary according to the PBE concept. The challenger's choice to escalate at that turn is made accordingly. In the second case, the supporter gives a commitment with probability $p = \frac{1}{3}$ and the challenger decides to test the waters with some fair probability $p \simeq 0.62069$. The defender resists with certainty if it has a firm commitment and does so with small probability $p = \frac{1}{6} \simeq 0.166667$ if it has no commitment. The equilibrium is therefore semi-separating at that stage. Nevertheless, the challenger continues with escalation in case of resistance. Is this what happened in July 1914?

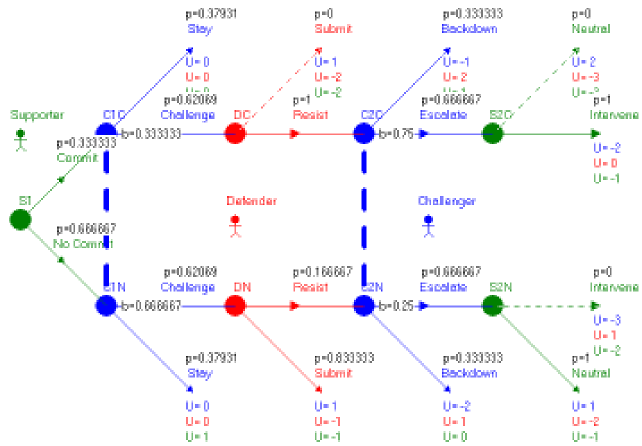

Figure 4.13: A Successful Alliance Deterrence

4.5.2 Perjury

This topic will be covered by the case study "Perjury Trap."

4.6 Homework

4.6.1

4.6.2

4.6.3

4.6.4

4.6.5