Lecture Notes on

Strategic Interactions: Games with Uncertain Preferences

Yakov Ben-Haim
Yitzhak Moda'i Chair in Technology and Economics
Faculty of Mechanical Engineering
Technion — Israel Institute of Technology
Haifa 32000 Israel
yakov@technion.ac.il
http://info-gap.com http://yakovbh.net.technion.ac.il/

Source material: Yakov Ben-Haim, and Keith W. Hipel, 2002, The graph model for conflict resolution with information-gap uncertainty in preferences, *Applied Mathematics and Computation,* 126: 319–340.

A Note to the Student: These lecture notes are not a substitute for the thorough study of books. These notes are no more than an aid in following the lectures.

Contents

[0]\lectures \risk\lectures\game02.tex 30.12.2020

1 Strategic Interactions

¶ **Strategic interaction:**

- Two or more protagonists.
- Each protagonist has 2 or more options.
- Each protagonist prioritizes his options based on:
  - His own preferences.
  - His knowledge of the preferences of the other protagonists.
- May be dynamic (multiple time steps) or static (one-shot).

¶ **Game theory:** versatile tool for studying strategic interactions.

¶ **Uncertainties:**

- Rules of interaction.
- Consequences (outcomes).
- Preferences of protagonists (including one's own).

¶ **We focus on uncertain preferences.**

¶ **Preliminary example:** The game of Chicken, considered again in section 4, p.12.

	State 1	State 2	State 3	State 4
Driver 1				
1. Swerve	N	N	Y	Y
Driver 2				
2. Swerve	N	Y	N	Y

Table 1: Four states in the game of chicken.

¶ **Rules of the game of 'chicken':**

1. Two car drivers race towards each other at high speed. Each driver has the choice of either swerving to avoid collision, or continuing to drive straight ahead.
2. Table 1 displays the game of chicken written in option form. The left hand side of the table lists each of the two decision makers in chicken, followed by the single option, or binary decision, called 'swerve', which each decision maker controls.
3. The four columns of Y's and N's represent the four possible states that could occur. A 'Y' indicates 'yes' the option is selected by the decision maker controlling it, while 'N' means the option is not taken.
4. For example, state 3 in Table 1: the first driver swerves while the second driver does not and thereby the second driver wins the game.

¶ **To swerve or not to swerve: how to decide?**

2 Info-Gap Model for Uncertain Transitive Preferences

2.1 Preference Vectors

¶ A decision maker's preferences are stated among specified **states.**

¶ J = number of different states.

¶ π = **preference vector.**
• π is a J-vector of integers.
• $\pi_m = r$ means that state m has preference rank r.
• Large r implies high preference for state m.
• $\pi_m = \pi_n$ means that states m and n have the same rank.
• Note that preferences are **transitive if:**
$\pi_i > \pi_j$ and $\pi_j > \pi_k$ then $\pi_i > \pi_k$.

¶ E.g.: suppose $J = 4$.
• $\pi = (4, 1, 3, 2)$ means:
$\pi_1 = 4$: State 1 has preference rank 4: most preferred.
$\pi_2 = 1$: State 2 has preference rank 1: least preferred.
$\pi_3 = 3$: State 3 has preference rank 3: next to least preferred.
$\pi_4 = 2$: State 4 has preference rank 2: 2nd most preferred.
• $\pi = (4, 1, 1, 1)$ means:
$\pi_1 = 4$: State 1 has preference rank 4: most preferred.
$\pi_2 = \pi_3 = \pi_4 = 1$: States 2, 3 and 4 have equal preference rank; less than state 1.

¶ Often one player will be **uncertain** about the
preference vector of another player.
• Each player's actions are influenced by his
knowledge of and **uncertainty about**
other player's preferences.
• Thus we are interested in **uncertain** preference vectors.
• We will formulate an **info-gap model** for
uncertain transitive preference vectors.

¶ A different info-gap model can be used to represent
uncertain preference vectors
where the preferences are possibly non-transitive.

¶ $\pi^{(0)}$ = known nominal preference vector.

2.2 An Info-Gap Model

¶ For two preference vectors, π and π', define the **distance** between them as the number of rank-changes by which they differ:

$$\text{dis}\,(\pi,\,\pi') = \sum_{j=1}^{J} \left|\pi_j - \pi'_j\right| \tag{1}$$

where $|x|$ is the ordinary absolute value of x, and π_i and π'_j are the jth elements of π and π'.

¶ Define the info-gap model as the set of preference vectors whose distance from the nominal preference vector $\pi^{(0)}$ is no more than h:

$$\mathcal{U}(h, \pi^{(0)}) = \left\{\pi :\ \text{dis}\,(\pi,\,\pi^{(0)}) \le h\right\}, \quad h = 0,\,1,\,2,\,\ldots \tag{2}$$

¶ $\mathcal{U}(h, \pi^{(0)})$:

- Is defined for any non-negative integer horizon of uncertainty h.
- Becomes more inclusive as h increases, so $\mathcal{U}(h, \pi^{(0)})$ is a family of **nested sets.**
- Contains only the nominal preference vector at $h = 0$, which is called "**contraction**".
- Contraction and nesting give h its meaning as an info-gap or horizon of uncertainty.

2.3 ‡ Preference-Change Operator and the Info-Gap Model

¶ This section develops formalism for the info-gap model of eq.(2), and is not needed for the rest of the lecture.

¶ e^i = column J-vector:
1 in ith position and 0 elsewhere.

¶ **Preference-change operator:**
For any preference vector π,
◦ $\pi + \mathrm{e}^i$ raises the rank of the ith state by 1.
◦ $\pi - \mathrm{e}^i$ lowers the rank of the ith state by 1.
◦ Technical proviso: these '+' and '−' operations do not raise a preference rank above J or lower it below 1.

¶ **Preference sets $\Pi^{(k)}$.**
- Define $\Pi^{(0)}$ as the set containing the nominal preference vector:

$$\Pi^{(0)} = \{\pi^{(0)}\} \tag{3}$$

- Recursively define the following sets of preference vectors:

$$\Pi^{(k)} = \left\{\pi \pm \mathrm{e}^i, \ \text{for all } \pi \in \Pi^{(k-1)}, \quad i = 1, \ldots, J\right\}, \quad k = 1, 2, \ldots \tag{4}$$

Thus $\Pi^{(k)}$ is the set of preference vectors
which differ from the nominal preference vector
by no more than k single preference changes.

¶ **Example.**
- The nominal preference vector is $\pi^{(0)} = (4, 1, 3, 2)$.
- $\Pi^{(0)}$ contains only $\pi^{(0)}$:

$$\Pi^{(0)} = \{(4, 1, 3, 2)\} \tag{5}$$

- $\Pi^{(1)}$ contains all 1-change variations on $\Pi^{(0)}$:

$$\begin{aligned}\Pi^{(1)} \;=\; & \{\underbrace{(\mathbf{4}, 1, 3, 2), (\mathbf{3}, 1, 3, 2)}_{\text{1st element}}, \underbrace{(4, \mathbf{2}, 3, 2), (4, \mathbf{1}, 3, 2)}_{\text{2nd element}} \\ & \underbrace{(4, 1, \mathbf{4}, 2), (4, 2, \mathbf{2}, 2)}_{\text{3rd element}}, \underbrace{(4, 1, 3, \mathbf{3}), (4, 2, 3, \mathbf{1})}_{\text{4th element}}\}\end{aligned} \tag{6}$$

- Etc.

¶ Now we define our **info-gap model of uncertainty** as:

$$\mathcal{U}(h, \pi^{(0)}) = \bigcup_{k=0}^{h} \Pi^{(k)}, \quad h = 0, 1, 2, \ldots \tag{7}$$

∘ $\mathcal{U}(h, \pi^{(0)})$ is the set of all preference vectors
which differ from the known nominal preference vector
by no more than h changes in rank.

• $\mathcal{U}(h, \pi^{(0)})$ is an info-gap model of uncertainty.

• h is the info-gap **horizon of uncertainty.**
h is a non-negative integer.
Meaning of h: max 'distance' of elements of $\mathcal{U}(h, \pi^{(0)})$
from the known nominal preference vector, $\pi^{(0)}$.

• The family of sets $\mathcal{U}(h, \pi^{(0)})$, $h = 0, 1, 2, \ldots$, are nested:

$$h < h' \quad \text{implies that} \quad \mathcal{U}(h, \pi^{(0)}) \subset \mathcal{U}(h', \pi^{(0)}) \tag{8}$$

3 Robustness Functions

¶ We now define several **robustness functions.**

¶ We consider N decision makers (protagonists).

- We are one of these decision makers.
- The preference vectors of the n decision makers are $\pi_1, \ldots, \pi_N$.
- $\pi_m^{(0)}$ is our best estimate of π_m, for $m = 1, \ldots, N$.
- We are uncertain about the preference vectors of some or all of the players. We may even be unsure about **our own** preferences.
- $\mathcal{U}_m(h, \pi_m^{(0)})$ is the info-gap model expressing our uncertainty about the preferences of the mth decision maker.

¶ We consider **three types of robustness functions:**

- Robustness of possible **resolutions of the conflict** as a whole (section 3.2, p.8).
- Robustness of a strategy of an **individual decision maker** (section 3.3, p.9)
- Robustness of a strategy with a **specified minimum acceptable ordinal utility** (section 3.4, p.11)

¶ Only sections 3.1 and 3.3 are needed for the example in section 4.

3.1 Equilibrium Solution Sets

¶ We will **not consider** the formal definition of:

— a game.

— an equilibrium solution of a game.

- Roughly: Suppose each player chooses an option.

  This is a **"solution".**

  An **equilibrium solution** is the situation where:

  — Each player has chosen an option.

  — No player has an incentive to change his choice once he learns the choices of the other players.
- Usually: at least one set of choices is an equilibrium resolution.
- Sometimes there will be **more than one** equilibrium resolution.
- For some game- and equilibrium-definitions we can prove that at least one equilibrium resolution exists.

¶ The set of all equilibrium solutions, based on the **nominal preferences,** $\pi_1^{(0)}, \ldots, \pi_N^{(0)}$ is the **nominal equilibrium solution set:** $\mathcal{Z}(\pi_1^{(0)}, \ldots, \pi_N^{(0)})$.

¶ The set of all equilibrium solutions, based on **arbitrary preferences,** $\pi_1, \ldots, \pi_N$ is the **equilibrium solution set:** $\mathcal{Z}(\pi_1, \ldots, \pi_N)$.

3.2 Robustness of a Conflict Resolution

¶ We will consider two concepts of robustness of the resolution of a conflict.

¶ The **overall robustness,** $\widehat{h}$, is the
greatest value of the uncertainty parameter h
for which the set of equilibrium solutions $\mathcal{Z}(\pi_1, \ldots, \pi_N)$
is the same as the nominal equilibrium set $\mathcal{Z}(\pi_1^{(0)}, \ldots, \pi_N^{(0)})$.
More precisely:

$$\widehat{h} = \max \left\{ h : \quad \mathcal{Z}(\pi_1, \ldots, \pi_N) = \mathcal{Z}(\pi_1^{(0)}, \ldots, \pi_N^{(0)}) \right. \\ \left. \text{for all} \quad \pi_m \in \mathcal{U}_m(h, \pi_m^{(0)}),\ m = 1, \ldots, N \right\} \tag{9}$$

Meaning of $\widehat{h}$:
○ Large $\widehat{h}$: the equilibrium solutions are robust to preference-uncertainty.
○ Small $\widehat{h}$: equilibrium solutions of the conflict are quite vulnerable
to the incomplete information available to the decision makers.

¶ A specialization of $\widehat{h}$ is where:
the preferences of some of the decision makers **are known.**
This is a **conditional robustness.**

¶ A **conditional robustness** is the
greatest value of the uncertainty parameter h
for which the set of equilibrium solutions $\mathcal{Z}(\pi_1, \ldots, \pi_N)$
is the same as the nominal equilibrium set $\mathcal{Z}(\pi_1^{(0)}, \ldots, \pi_N^{(0)})$.
for fixed values of the preferences of some of the decision makers.

¶ Let $\mathcal{I}$ be a set of indices of some but not all of the decision makers.
• The robustness conditioned on fixed preferences for decision makers indexed in $\mathcal{I}$ is:

$$\widehat{h}(\pi_m,\ m \in \mathcal{I}) = \max \left\{ h : \quad \mathcal{Z}(\pi_1, \ldots, \pi_N) = \mathcal{Z}(\pi_1^{(0)}, \ldots, \pi_N^{(0)}) \right. \\ \left. \text{for all} \quad \pi_n \in \mathcal{U}_n(h, \pi_n^{(0)}),\ n \notin \mathcal{I} \right\} \tag{10}$$

¶ We see that, if $\mathcal{I} = \emptyset$ then:

$$\widehat{h} = \widehat{h}(\pi_m,\ m \in \mathcal{I}) \tag{11}$$

3.3 Robustness of a Strategy

¶ **Strategies.**

1. We consider a situation in which each decision maker may use a **strategy** by which he influences his own as well as his opponents' reachable states.
2. In a game-theoretic setting, a **strategy** σ at each stage of the game is simply a choice from the set of options available to that decision maker.
3. We will define two robustnesses associated with the choice of a strategy.

¶ **Solution sets.**

1. The equilibrium solution set $\mathcal{Z}(\pi_1, \ldots, \pi_N)$ is defined as before.
2. Let $\mathcal{Z}_m(\pi_1, \ldots, \pi_N|\sigma)$ be the subset of $\mathcal{Z}(\pi_1, \ldots, \pi_N)$ whose elements all entail the implementation of strategy σ by decision maker m.

¶ The **overall robustness** of strategy σ
is the greatest value of the uncertainty parameter h
for which the **equilibrium solution** subset $\mathcal{Z}_m(\pi_1, \ldots, \pi_N|\sigma)$
is the same as
the **nominal equilibrium solution** subset, $\mathcal{Z}(\pi_1^{(0)}, \ldots, \pi_N^{(0)})$.
if m implements σ:

$$\gamma_m(\sigma) = \max \left\{ h : \quad \mathcal{Z}_m(\pi_1, \ldots, \pi_N|\sigma) = \mathcal{Z}_m(\pi_1^{(0)}, \ldots, \pi_N^{(0)}|\sigma) \right. \\ \left. \text{for all} \quad \pi_n \in \mathcal{U}_n(h, \pi_n^{(0)}),\ n = 1, \ldots, N \right\} \tag{12}$$

$\gamma_m(\sigma)$ is analogous to the overall robustness, $\widehat{h}$, eq.(9), p.8,
where the additional constraint is imposed that
decision maker m implements strategy σ.

¶ Now consider a **conditional robustness** based on $\gamma_m(\sigma)$:

$$\gamma_m(\pi_m,\ m \in \mathcal{I}|\sigma) = \max \left\{ h : \quad \mathcal{Z}_m(\pi_1, \ldots, \pi_N|\sigma) = \mathcal{Z}(\pi_1^{(0)}, \ldots, \pi_N^{(0)}|\sigma) \right. \\ \left. \text{for all} \quad \pi_n \in \mathcal{U}_n(h, \pi_n^{(0)}),\ n \notin \mathcal{I} \right\} \tag{13}$$

$\gamma_m(\sigma)$ is analogous to the conditional robustness $\widehat{h}(\pi_m,\ m \in \mathcal{I})$, eq.(10), p.8.
We have imposed upon $\gamma_m(\sigma)$
the further condition that particular preferences are known and fixed.

¶ The use of $\gamma_m(\pi_m, m \in \mathcal{I}|\sigma)$:

1. Evaluating and choosing between alternative strategies.
2. Suppose that :
   (a) decision maker m is contemplating two alternative strategies, σ_1 and σ_2.
   (b) Their robustnesses are ranked as:
$$\gamma_m(\pi_m, m \in \mathcal{I}|\sigma_1) \gg \gamma_m(\pi_m, m \in \mathcal{I}|\sigma_2) \tag{14}$$
3. The decision maker may tend to prefer σ_1 over σ_2.
4. This conclusion is reasonable even if the nominal preference of decision maker m for σ_1 is less than his nominal preference for σ_2.
5. In other words, the analysis of the robustness to uncertainty can alter a decision maker's preferences.

¶ **When is the robust preference in eq.(14) not reasonable?**

- If the nominal outcome of the strategy σ_1 is unacceptable.
- This motivates the next concept of robustness.

3.4 Robustness of a Strategy with Minimum Acceptable Utility

¶ **Critical Utility,** u_c.

1. A decision maker may be unable to obtain the state that he most prefers.
2. Or, attainment of that state may be very vulnerable to uncertainty.
3. Thus, the decision maker may wish to identify a **critical utility,** u_c which is an integer corresponding to the least ordinal preference he is willing to accept.
4. If he insists on the optimal, then $u_c = J$; if he is more modest, he will choose a value less than J for u_c.

¶ When decision maker m contemplates a strategy σ,
he may wish to know how robust the implementation of that strategy is,
to uncertainty in the preferences of the other decision makers.
◦ Also, he may wish to know
how the robustness of a strategy varies with u_c,
the least-acceptable ordinal utility.

¶ **Minimal utility function** $u_m(\pi_1, \ldots, \pi_N|\sigma)$:
The least ordinal utility that the decision maker could achieve,
at equilibrium,
if he implements strategy σ and if
the preferences of the decision makers are $\pi_1, \ldots, \pi_N$.

¶ The robustness of strategy σ for decision maker m,
when he demands ordinal utility no less than u_c, is:

$$\zeta_m(u_c, \sigma) = \max \left\{ h : \quad u_m(\pi_1, \ldots, \pi_N|\sigma) \geq u_c, \quad \text{for all} \quad \pi \in \mathcal{U}_m(h, \pi_m^{(0)}),\ m = 1, \ldots, N \right\} \tag{15}$$

¶ The use of the robustness $\zeta_m(u_c, \sigma)$ is in evaluating and choosing between alternative strategies, similar to the use of $\gamma_m(\pi_m,\ m \in \mathcal{I}|\sigma)$. In other words, $\zeta_m(u_c, \sigma)$ can be used to evaluate and revise a decision maker's preferences in light of the uncertainty of his knowledge.

¶ Another important property of this robustness is the trade-off between demanded utility and robustness. When u_c is very large, meaning that the decision maker is demanding a very favorable outcome, then the robustness of any strategy σ will be low. That is, $\zeta_m(u_c, \sigma)$ decreases monotonically as u_c increases.

4 The Game of 'Chicken'

1. We will perform an uncertainty analysis of the game of 'chicken'.
2. We will use the conditional robustness of a strategy, defined in eq.(13) on p. 9.
3. We will demonstrate that the analysis of uncertainty can alter a decision maker's priorities, and can lead to decisions which are different from those made in the absence of uncertainty.

	State 1	State 2	State 3	State 4
Driver 1				
1. Swerve	N	N	Y	Y
Driver 2				
2. Swerve	N	Y	N	Y

Table 2: Four states in the game of chicken.

¶ Rules of the game of 'chicken':

1. Two car drivers race towards each other at high speed. Each driver has the choice of either swerving to avoid collision, or continuing to drive straight ahead.
2. Table 2 displays the game of chicken written in option form. The left hand side of the table lists each of the two decision makers in chicken, followed by the single option, or binary decision, called 'swerve', which each decision maker controls.
3. The four columns of Y's and N's represent the four possible states that could occur. A 'Y' indicates 'yes' the option is selected by the decision maker controlling it, while 'N' means the option is not taken.
4. For example, state 3 in Table 2: the first driver swerves while the second driver does not and thereby the second driver wins the game.

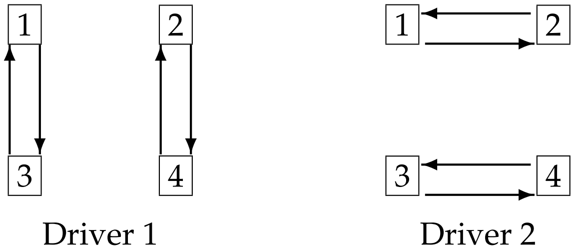

Figure 1: Unilateral moves by the two drivers in the game of chicken.

¶ **Directed Graph of the game of chicken:**

1. Fig. 1 portrays the directed graph of unilateral moves in one step for each of the decision makers.
2. When the second driver remains fixed on choosing N for his strategy, the first driver can unilaterally cause the game to move from state 1 to 3 by changing his selected option from N to Y.
3. If driver 1 is allowed to go back onto the road after swerving, he also controls the movement from state 3 to 1.
4. Driver 1 also is in charge of the movement between states 2 and 4 when the second driver has a fixed strategy of Y.
5. Following similar arguments, driver 2 controls the unilateral movements between states 1 and 2 and well as states 3 and 4.

¶ **Our side of the game:**

1. We are driver 1.
2. We adopt the preference vector $\pi_1 = (1, 4, 2, 3)$.
3. This means:
   (a) Our greatest preference is for state 2, in which our opponent swerves but we do not.
   (b) Our next preference is for state 4 (both swerve), then for state 3 (only I swerve) and our least preferred outcome is state 1 (collision).
   (c) This is a reasonable and rather humane preference vector. Our first preference is for state 2, in which only our opponent swerves and is disgraced. However, by giving our next preference to state 4, we are emphasizing an 'amicable' outcome, in which both drivers swerve, both survive, and both are 'disgraced' to the same degree.

¶ **The opponent:**

1. Our opponent's preference vector is $\pi_2 = (1, 1, 4, 1)$, which is quite different from ours.
2. His only preference is for disgracing us by forcing us to swerve; he is indifferent to any other outcome.

¶ **Ignore uncertainty for the moment.**

¶ We will adopt **Nash stability:**

1. A state is **stable** for a given decision maker if he has no incentive for moving from that state.
2. That is, state m is Nash stable for decision maker i if i's preference for state m is no less than his preference for any other accessible state for fixed strategy choices of the other decision makers.
3. State m is stable for decision maker i if:

$$\pi_{m,i} \geq \pi_{k,i} \tag{16}$$

for all states k
in which the strategy choices of the other decision makers
are the same as in state m.

¶ **Equilibrium:**
A state is an **equilibrium** if it is stable for both decision makers.

State	decision maker 1	decision maker 2	Outcome
1	unstable	stable	not equilibrium
2	stable	stable	equilibrium
3	stable	stable	equilibrium
4	unstable	unstable	not equilibrium

Table 3: Nash stability for the game of chicken, with the preference vectors $\pi_1 = (1, 4, 2, 3)$ and $\pi_2 = (1, 1, 4, 1)$.

¶ **Stable and equilibrium states for 'chicken',** table 3:

1. State 3 is Nash stable for driver 1 because he prefers state 3 more than state 1 and from 3 he can only move to 1.
2. Notice for both states 3 and 1 driver 2 has a fixed strategy of N for not swerving. From driver 2's point of view, state 3 is Nash stable for him because he prefers state 3 over state 4, where driver 1 has a fixed strategy of Y. Since state 3 is Nash stable for both drivers, it constitutes a Nash equilibrium.
3. States 2 (only he swerves) and 3 (only I swerve) are both equilibrium outcomes: neither decision maker has an incentive to move from either of these states, according to the specified preference vectors and Nash stability definition.

¶ **Analysis of uncertainty** from 'our' viewpoint: driver 1.
• We employ the robustness of eq.(13) on p. 9.
• We know our own preference vector precisely: $\pi_1 = (1, 4, 2, 3)$.
• However, our knowledge of our opponent's preferences is uncertain.
• Our nominal preference vector for our opponent is $\pi_2^{(0)} = (1, 1, 4, 1)$.

¶ **Nominal equilibrium solutions, table 3:**
◦ We see that the set of equilibrium states, based on the nominal preference vectors, contains states 2 (only I don't swerve) and 3 (only I swerve):

$$\mathcal{Z}(\pi_1, \pi_2^{(0)}) = \{2, 3\} \tag{17}$$

◦ So, the set of equilibrium solutions constrained to entail option 'No' by driver 1 is the subset of $\mathcal{Z}(\pi_1, \pi_2^{(0)})$ which contains state 2:

$$\mathcal{Z}_1(\pi_1, \pi_2^{(0)} | \sigma = \text{No}) = \{2\} \tag{18}$$

◦ Likewise, the set of equilibrium solutions constrained to entail option 'Yes' by driver 1 is the subset of $\mathcal{Z}(\pi_1, \pi_2^{(0)})$ which contains state 3:

$$\mathcal{Z}_1(\pi_1, \pi_2^{(0)} | \sigma = \text{Yes}) = \{3\} \tag{19}$$

Level of Uncertainty (h)	Equilibrium Sets $\mathcal{Z}(\pi_1, \pi_2)$ $\pi_2 \in \mathcal{U}(h, \pi_2^{(0)})$	Constrained Eq. Sets $\mathcal{Z}_1(\pi_1,\ \pi_2 \vert \sigma = \text{Yes})$ $\pi_2 \in \mathcal{U}(h, \pi_2^{(0)})$	Constrained Eq. Sets $\mathcal{Z}_1(\pi_1,\ \pi_2 \vert \sigma = \text{No})$ $\pi_2 \in \mathcal{U}(h, \pi_2^{(0)})$
0	{2, 3}	{3}	{2}
1	{2, 3}, {3}	{3}	{2}, ∅
2	{2, 3}, {3}	{3}	{2}, ∅
3	{2, 3}, {3}	{3}	{2}, ∅
4	{2, 3}, {2}, {3}	{3}, ∅	{2}, ∅

Table 4: Equilibrium sets with increasing uncertainty with preference vectors $\pi_1 = (1, 4, 2, 3)$ and $\pi_2^{(0)} = (1, 1, 4, 1)$.

¶ **Equilibrium solutions with uncertainty, table 4:**

1. $h = 0$:
   (a) No uncertainty in the preference vector for driver 2.
   (b) The set of equilibrium solutions contains states 2 and 3, as we expect from Table 3.
   (c) The constrained equilibrium sets contain states 2 or 3, for the two different options considered by driver 1, as explained in connection with eqs.(18) and (19).
2. $h = 1$:
   (a) Uncertainty of no more than one step away from $\pi_2^{(0)}$.
   (b) More than a single equilibrium solution set exists: the previous set as well as the set containing only state 3.
   (c) Now the constrained solution set for option $\sigma =$ 'Yes' contains only state 3.
   (d) However, for the option $\sigma =$ 'No', one constrained equilibrium set is $\{2\}$ while the other constrained set is empty, implying that the desired option does not lead to equilibrium with some preference vector for driver 2 at uncertainty $h = 1$.
3. $h = 2$ and $h = 3$: same as $h = 1$.
4. $h = 4$:
   (a) One of the constrained equilibrium solutions sets for option $\sigma =$ 'Yes' is empty.
   (b) This means that a preference vector for driver 2 at uncertainty $h = 4$ exists which does not allow equilibrium if driver 1 chooses this option.

¶ **Conclusion from Table 4:**

1. If driver 1 chooses option $\sigma =$ 'Yes', then he will always reach an equilibrium, for any preference vector for driver 2 up to and including uncertainty level $h = 3$.
2. On the other hand, one cannot guarantee that the option $\sigma =$ 'No' will be an equilibrium point at uncertainty level $h = 1$.
3. In short, the conclusion is that the conditional robustness, eq.(13), for the two options available to driver 1 are:

$$\gamma_1(\pi_1|\sigma = \text{Yes}) = 3 \tag{20}$$
$$\gamma_1(\pi_1|\sigma = \text{No}) = 0 \tag{21}$$

¶ **Summary:**

1. Option 'Yes' is much more robust than option 'No', with respect to driver 1's uncertainty about his opponent's preferences.
2. On the one hand, state 2 (entailing option 'No') is preferred by driver 1 over state 3 (entailing option 'Yes'). Recall that driver 1's preference vector is: $\pi_1 = (1, 4, 2, 3)$.
3. However, the analysis of uncertainty has shown that the achievement of equilibrium for state 2 is much more fragile to uncertainty than achievement of equilibrium for state 3.
4. Consequently, driver 1 may well select option 'Yes' (state 3) rather than 'No' (state 2).
5. What we see, in short, is that the analysis of uncertainty can fundamentally alter a decision maker's preferences.