Game States

Aviad Heifetz*

Very Preliminary, April 2006

Abstract

We provide an inductive definition of *game-positions,* which generalize the notion of nodes in perfect-information games with no infinite paths. *Game-states* are defined inductively as well. Essentially, a game-state consists of a specification of the game-states that result from players' actions, and the players' types, specifying their beliefs on game-states. In a *dynamic equilibrium* each type chooses an action to maximize her expected payoff according to her belief, given the actions of other types.

Dynamic equilibrium generalizes the notion of Bayes-Nash equilibrium to a dynamic setting in which a common prior or Bayesian updating are not imposed a priori. Various formulations of beliefs enable further generalizations, expressing unawareness or imperfect recall.

1 Short Introduction

In the classical definition of extensive-form games with imperfect information, the notion of an *information set* is half-baked. On one hand, it is meant to specify objective circumstances: part of the development of the game this far has been (physically) concealed from the player. On the other hand, it specifies a subjective state of mind of the player. In the latter dimension, it is restrictive in at least three senses.

First, it does not allow for the possibility of mistakes on the part of the player, who always considers as possible the development of the game which actually took place.

Second, it specifies only a qualitative belief – these and those nodes are possible. It leaves the quantitative part (e.g. probabilistic beliefs on the nodes) to be specified as part of the equilibrium notion (e.g. sequential equilibrium).

*The Economics and Management Department, The Open University of Israel, avaidhe@openu.ac.il.

Third, the uncertainty expressed by the information set is – strictly speaking – about the realized path of actions in the game, and not about the other players' state of mind. To capture the latter kind of uncertainty, one has to have, in a somewhat roundabout way, a chance move at the beginning of the game tree leading to many combinations of types, which would differ from one another by their beliefs following some history of the game. This description is suitable for a strategic interaction which can be fully described in a Harsanyi-consistent way before the game starts – the preliminary chance move specifies the common prior over the types, and these types remain the same along the entire game. The description is not suitable when states of mind of the player are not compatible with a common prior.

The definitions that follow aim at overcoming these difficulties. The cognitive structure – types in *game-states* – is specified on top and separately of the initial, physical description of *game-positions.* Game-positions are defined inductively, and generalize the notion of nodes in perfect-information games with no infinite paths. The definition of game-states is inductive as well.

We first pursue a definition of game-states in which types are defined by probabilistic beliefs on game-states. When each type chooses an action so as to maximize her expected payoff given the actions chosen by other types, we get a *dynamic equilibrium.*

Does the cognitive structure, as specified by the game states, form part of the definition of a dynamic equilibrium, or is this cognitive structure well-defined independently, before the equilibrium actions are specified? The former is the case if types are further assumed – implicitly or explicitly – to be Bayesian updaters, because beliefs about game-states are then related to the actions chosen by types in preceding game-states. The latter is the case if Bayesian updating is assumed away, or if it holds vacuously. For example, Bayesian updating is irrelevant when the game is essentially a one-shot simultaneous-move Bayesian game. Bayesian updating is irrelevant also if it is commonly believed that each player is called to play only once, and in this occasion the player forms beliefs only about the current and future game-states. We will analyze such a particular, non-trivial example.

In the sequel,

- we regress to show that if, alternatively, the state of mind of a type is characterized by a possibility set rather than by a probabilistic belief, then information sets (in the classical definition of imperfect-information extensive-form games) can be formulated as a particular case. Some kinds of imperfect recall are allowed as well, as long as no path contains two distinct game-states in which one of the players has the same type.

- we progress to show how to describe *games with unawareness* when types are generalized so as to capture states of mind of partial and mutual unawareness.

- We generalize the definition of types further, so as to capture imperfect recall, where – as in the absent-minded paradox – there exists a path in which a player has the same type in two distinct game-states along the path.

2 Game-positions

Game-positions are defined inductively. Their definition constitutes a generalization of the definition of nodes in perfect-information extensive-form games with no infinite paths (when, as in Osborne and Rubinstein 1994 and Osborne 2005, both chance moves and simultaneous moves of several players are allowed in each node).

I is a set of players.

The set of payoffs to a player $i \in I$ is a subset $L_i \subseteq \mathbb{R}$, equipped with its Borel σ-field.

We are now going to define, inductively, spaces of *game-positions*.

The measurable space of payoff vectors $L = \prod_{i \in I} L_i$ is a space of game-positions, called the space of *leaf* game-positions (or terminal game-positions).

Let $\tilde{P}$ be a measurable space of game-positions, and A_i a measurable space of *actions* for each player $i \in I_0 \subseteq I$. Denote by $A = \prod_{i \in I_0} A_i$ a space of joint actions of the players in I_0. A space P, consisting of measurable functions

$$p : A \to \Delta\left(\tilde{P}\right) \tag{1}$$

from joint actions A to probability distributions over game-positions in $\tilde{P}$, is a space game-positions. [1] ,[2]

[1] If (X, Σ) is a measurable space, $\Delta(X)$ is a measurable space with the σ-field generated by the sets

$$B^p(E) \equiv \{\mu \in \Delta(X) : \mu(E) \geq p\}$$

where $E \in \Sigma$ and $p \in [0,1]$ is rational.

[2] $P \subset \Delta\left(\tilde{P}\right)^A$ is endowed with the σ-algebra it inherits from the product measurable space $\Delta\left(\tilde{P}\right)^A$ (generated by the cylinders of the form $\prod_{a \in A} Y_a$, where Y_a is measurable in $\Delta\left(\tilde{P}\right)$ for all $a \in A$ and $Y_a = \Delta\left(\tilde{P}\right)$ for all but finitely many $a \in A$.)

If $\mathcal{P}$ is a set of spaces of game-positions, their union is a space of game positions.[3]

This concludes the inductive definition of the class of game positions.

If, in the inductive definition above, we replace (1) by

$$p : A \to \tilde{P} \tag{2}$$

we get an inductive definition of the sub-class of *deterministic game-positions.* The other game positions are called *non-deterministic game-positions.*

In the general class of game-positions, chance moves are represented by non-deterministic game-positions in which $I_0 = \emptyset$ so that $A = \{\emptyset\}$, and $p(\emptyset)$ is a probability distribution over $\tilde{P}$.

2.1 Paths of game-positions

A *path* or *history* is a sequence $p_1, p_2, \ldots$ of deterministic game positions such that p_{k+1} is in the range of p_k for all $k \geq 1$

It is immediate to verify, by induction on the structure of spaces of game-positions, that there exist no infinite paths: Paths starting from the space of leaf game-positions have length 1; if P is a $\tilde{P}$-based space of game positions, and by the induction hypothesis all the paths starting from game-positions $\tilde{p} \in \tilde{P}$ are finite, then so are paths starting from game-positions $p \in P$; and if, by the induction hypothesis, only finite paths start from game positions in all the spaces of game-positions in some set $\mathcal{P}$ of spaces of games-positions, then so is the case for the game positions in the union of the spaces in $\mathcal{P}$.

Notice that the paths that start from a game-position p span a directed graph which is not necessarily a tree. This is because distinct paths that start in p may end in the same game-position p'. In this sense, game-positions generalize the notion of nodes in extensive-form games.

The notion of a path can be extended to contain also non-deterministic game-positions p, provided that the definition of the probability measure $p(a) \in \Delta\left(\tilde{P}\right)$, $a \in A$ is augmented by a notion of a support – a set $\tilde{P}' \subseteq \tilde{P}$ such that $p(a)\left(\tilde{P}'\right) = 1$. We then require, in the extended definition of a path, that p immediately precedes $\tilde{p}$ in the path only if $\tilde{p}$ is in the support of $p(a)$ for some $a \in A$.

The class of game-positions p which start no distinct paths ending in the same game-position p', is called the the class of *classical game-positions.* These are isomorphic to the class of nodes in perfect-information extensive-form games.

[3] The union of $\mathcal{P}$ is endowed with the σ-algebra generated by the union of the σ-algebras of the members of $\mathcal{P}$.

3 Game-states

A game-position was defined by specifying what game position (or distribution of game positions) follows as a result of players' actions.

A game-state specifies not only the consequences of players' actions, but also the players' beliefs about these consequences and about the other players' beliefs. This is accomplished by specifying what game-state (or distribution of game states) results from the players' actions, and what is the type of each player. A type of a player specifies the belief of the player on actions' consequences and on other players' types.

We proceed with the formal, inductive definition of *game-states*.

Leaf game-positions are game-states.

Let $\tilde{\Omega}$ be a measurable space of game-states, and A_i a measurable space of actions for each player $i \in I_0 \subseteq I$. Denote by $A = \prod_{i \in I_0} A_i$ a space of joint actions of the players in I_0, and by K a space of measurable functions

$$k : A \to \Delta\left(\tilde{\Omega}\right) \tag{3}$$

from joint actions A to probability distributions over game-states in $\tilde{\Omega}$. Let T_i be a measurable space of *types* for each player $i \in I_0$, and

$$b_i : T_i \to \Delta\left(K \times T_{-i}\right) \tag{4}$$

a measurable function that specifies a probabilistic belief of each type $t_i \in T_i$ on K and the other players' types $T_{-i} = \prod_{j \in I_0 \setminus i} T_j$. Then

$$\Omega = K \times \prod_{i \in I_0} T_i$$

is a space of $\tilde{\Omega}$-based game-states.

Any union of a set of spaces of game-states is a space of game-states.

This concludes the inductive definition of spaces of game-states.

If, in the above inductive definition we replace (3) by

$$k : A \to \tilde{\Omega} \tag{5}$$

we get the sub-class of *deterministic game-states*.

3.1 Identifying the game-positions in game-states

We now define, inductively, a measurable map Π that specifies the game-position in each game-state. The idea is simple: If the game-state is a leaf game-position, then the game-position is already identified. If, inductively, the game-positions of the game-states in $\tilde{\Omega}$ have already been identified, then game-states in Ω – the actions in which lead to game-states in $\tilde{\Omega}$ – can now be associated with game-positions in which the same actions lead to the corresponding game-positions identified for $\tilde{\Omega}$.

Here is the formal inductive definition of the map Π:

Π is the identity map on the space of game-states which are leaf game-positions.

If Ω is a space of $\tilde{\Omega}$-based game-states, then for each $\left(k, (t_i)_{i\in I_0}\right) \in \Omega$, the game-position $\Pi\left(\left(k, (t_i)_{i\in I_0}\right)\right)$ is defined by

$$\Pi\left(\left(k, (t_i)_{i\in I_0}\right)\right)(a) = \Pi\left(k(a)\right), \qquad a \in A$$

(If case $k(a)$ is a probability measure in $\Delta\left(\tilde{\Omega}\right)$, then $\Pi\left(k(a)\right)$ is a probability measure on $\Pi\left(\tilde{\Omega}\right)$ defined by $\Pi\left(k(a)\right)(E) = k(a)\left(\Pi^{-1}(E)\right)$.)

If the game-positions in $\Pi(\Omega)$ are all classical, we say that Ω is a space of *classical game-states.*

3.2 Strategies

Let T_i be the set of types of player $i \in I_0$ in the space of game-states $\Omega = K \times \prod_{i\in I_0} T_i$. A *pure strategy* of player i is a measurable function

$$\sigma_i : T_i \to A_i$$

A *behavioral strategy* of player i is a measurable function

$$\sigma_i : T_i \to \Delta(A_i)$$

If Ω is a space of game-states which is a union of the disjoint spaces $\left\{\Omega^\alpha = K^\alpha \times \prod_{i\in I_0^\alpha} T_i^\alpha\right\}$ then a pure/behavioral strategy for player i is defined in Ω by defining it in each of the spaces Ω_α.

3.3 Paths of game-states

In the discussion that follows we shall restrict attention to pure strategies and deterministic game-states. This is done only for expositional simplicity, to avoid multiple integrals in the definitions. The definitions for the general case appear in an appendix.

Once the players' strategies σ_i, $i \in I_0$ are defined in Ω, we can define for every game-state $\omega = \left(k, (t_i)_{i\in I_0}\right) \in \Omega$ (i.e., a game-state which is not a leaf) the *consecutive* or *successor* game-state $\tilde{\omega} \in \tilde{\Omega}$ by

$$\tilde{\omega} = k\left(\sigma_i\left(t_i\right)_{i\in I_0}\right)$$

If for all such $\omega = \left(k, (t_i)_{i\in I_0}\right) \in \Omega$ the successor games $\tilde{\omega}$ are leaf games, the expected payoff to type t_i is

$$U_i\left(t_i; (\sigma_j)_{j\in I_0}\right) = \int_{K\times T_{-i}} k\left(\sigma_i\left(t_i\right), \sigma_j\left(t_j\right)_{j\neq i}\right)_i \mathrm{d}b_i\left(t_i\right)$$

In such a case, replacing each $\omega = \left(k, (t_i)_{i\in I_0}\right) \in \Omega$ by the leaf game-position

$$\left(U_i\left(t_i; (\sigma_j)_{j\in I_0}\right)\right)_{i\in I_0}$$

is called the *folding* of Ω with the strategies $(\sigma_j)_{j\in I}$. After the folding, all the game-states in Ω become leafs.

Let $\Omega_1, \ldots, \Omega_{n+1}$ be a finite sequence of spaces of game-states, where $\Omega_{n+1} = L$ and Ω_m is a space of Ω_{m+1}-based game-states for $m = 1, \ldots, n$. Let strategies be specified for all players in $\Omega_1, \ldots, \Omega_n$. Then the process of folding can be applied inductively to Ω_n, then to Ω_{n-1}, etc. At the end of the process, Ω_1 will be replaced by a space of leaf game-positions $(\ell_i)_{i\in I}$. The number $\ell_i \in \mathbb{R}$ is called the payoff of player i in this folding.

More generally, by induction on the structure of game-state spaces, every game-state space Ω can be folded, provided that strategies have been specified in all game-state spaces encountered up to (and including) the definition of Ω.

A *path* or *history* of game-states is a sequence $\omega_1 = \left(k_1, (t_i)_{i\in I_0^1}\right), \omega_2 = \left(k_2, (t_i)_{i\in I_0^2}\right), \ldots$ of deterministic game-states such that ω_{m+1} is in the range of k_m for all $m \geq 1$. One can verify, by induction on the structure of spaces of game-positions, that there exist no infinite paths of game-states. If strategies σ_i are defined for all the states in the path $\omega_1, \ldots \omega_{m+1}$, the path is called a *realized path by the strategies* σ_i if $\omega_{m+1} = k_m\left(\sigma_i\left(t_i\right)_{i\in I_0^m}\right)$ for all $m = 1, \ldots, n$.

3.4 Dynamic equilibrium

A game-state space Ω is called path-closed if it contains all the game-states that belong to paths starting at game-states in Ω. A tuple of pure strategies $(\sigma_i)_{i \in I}$ is a *dynamic equilibrium* at a path-close Ω if when Ω is folded with these strategies, no player at no game-state would get a higher payoff in the folding that would result by unilaterally altering her chosen action at that game-state.

The set of dynamic equilibria can be found by the usual process of backward induction, defined by induction on the structure of the game-state spaces of which Ω is a union.

4 Examples

4.1 Players who doubt others' faultlessness

The following example is a variant of an example from Heifetz and Pauzner (2005). It describes a finite-horizon overlapping generations model with fiat money.

There are 3 players: A grandmother (player 3), her daughter (player 2) and the granddaughter (player 1). Each of them is endowed with a physical good, consumption of which ("a_k" - autarky, $k = 1, 2, 3$) yields the player a payoff of 1. But the grandmother values the endowment of her daughter at 8, and similarly the daughter values the endowment of the granddaughter at 8. The grandmother has the option to forego ("f_3") the consumption of her endowment in exchange for a piece of paper called "dollar". If, afterwards, the daughter opts for autarky (a_2), the grandmother remains with a worthless piece of paper and her payoff is 0. If, alternatively, the daughter forgoes (f_2) her endowment in exchange for the grandmother's dollar, the grandmother's payoff is 8. If, then, the granddaughter forgoes (f_1) her endowment in exchange for the dollar, her payoff is 0 (because the granddaughter has no descendant) while the daughter gets a payoff of 8. If, alternatively, the granddaughter opts for autarky (a_1), she gets a payoff of 1, and the daughter remains with the worthless dollar and a payoff of 0. If any of the players chooses autarky, then all the subsequent players can only consume their own endowment, yielding them a payoff of 1.

The granddaughter has two types. With type '1', she believes that her payoff function is as described above. With type '0' ("confused"), she believes that her payoffs are preceded with a minus sign.

The daughter has four types: 11,10,01,00. If the first (i.e. right-most) digit is 1, the type believes that the granddaughter's payoffs are as described, and if the first digit is

0, the type believes that the payoffs of the granddaughter are preceded with a minus sign. Similarly, If the second digit is 1, the daughter believes that her own payoffs are as described above, and otherwise she believes her payoffs are preceded by a minus sign.

What probability does a type $t_2 = (t_2^2 t_2^1)$ of the daughter assign to a type $t_1 = (t_1^1)$? The probability is $\varepsilon = 0.1$ if $t_2^1 \neq t_1^1$, and $1 - \varepsilon = 0.9$ if $t_2^1 = t_1^1$. That is, player 2 – the daughter – assigns a high probability that her belief about the granddaughter's payoffs coincide with the granddaughter's belief about these payoffs, and the complementary probability that these beliefs are distinct.

The grandmother has 8 types – $\{0,1\}^3$. If the k-th digit in the type is "1", the type believes that the payoffs of player k are as described above, and if the k-th digit in the type is "0", the type believes that the payoffs of player k are preceded by a minus sign.

Type $t_3 = (t_3^3 t_3^2 t_3^1)$ of the grandmother assigns to the type $t_2 = (t_2^2 t_2^1)$ of the daughter the probability $(1-\varepsilon)^d \varepsilon^{2-d}$, where d is the number of indices on which their types "agree":

$$d = \# \left\{ k \in \{1,2\} : t_3^k = t_2^k \right\}$$

One can now compute, by backward induction, the optimal action of each type, assuming expected payoff maximization.

Type 1 of the granddaughter obviously opts for autarky, while type 0 prefers to forego her endowment (which she believes to entail her a payoff of -1) if she gets the option to do so.

If the daughter is of types 10 or 01, she prefers to forego her endowment: Type $t_2^2 t_2^1 = 10$ believes with a high probability that the granddaughter is of type 0 and will like to trade her endowment for a dollar; type $t_2^2 t_2^1 = 01$ assigns a high probability that the granddaughter is of type 1 and wouldn't materialize her option to trade, in which case the worthless dollar is better for the daughter than her endowment (which the daughter believes to entail her a payoff of -1).

The type 111 of the grandmother foregoes her endowment for a dollar: She assigns probability

$$0.9 \times 0.1 + 0.1 \times 0.9 = 0.18$$

to the event that the daughter is of the type 10 or 01 and would forego her endowment for the dollar, an exchange that would yield the grandmother an expected payoff of

$$0.18 \times 8 = 1.44$$

which is larger than the payoff of 1 that the grandmother can guarantee to herself by autarky.

In fact, all the types $t_3 = (t_3^3 t_3^2 t_3^1)$ of the grandmother in which $t_3^3 = 1$ (the types who believe that the payoff of the grandmother are as described and *not* preceded by a minus sign) forego the endowment for a dollar: All these types assign to the event that the daughter is of the type 10 or 01 either the probability 0.18 (in case $t_3^2 t_3^1 = 11$ as analyzed above, or in case $t_3^2 t_3^1 = 00$), or the probability

$$0.9 \times 0.9 + 0.1 \times 0.1 = 0.82$$

(in case $t_3^2 t_3^1 = 10$ or $t_3^2 t_3^1 = 01$). For these latter types, the expected payoff to the grandmother from foregoing her endowment for a dollar is

$$0.82 \times 8 = 6.56$$

which is higher than the payoff of 1 she gets from autarky.

Consequently, all the types $t_3 = (t_3^3 t_3^2 t_3^1)$ of the grandmother in which $t_3^3 = 0$ prefer autarky, because these types believe that all the payoffs of the grandmother are preceded by a minus sign.

Notice that the types in this example do not emerge from a common prior. Even under an interpretation with which there exists some "objective" probability by which the types are actually drawn, the beliefs of the types are not the posteriors of this "objective" probability, and each type is firm in her beliefs about her payoff and about the types of the player(s) that play after her. That's why these types cannot be described as the result of a chance move in the classical framework of extensive-form games.

Notice also that in this example, a player does not have beliefs about the types of the player(s) that precede her in the game. The player understands that different actions are optimal for different types of the preceding player or players. However, since the player does not have a belief on these types to start with, observing the preceding players' actions does not cause any update in the beliefs of the player about subsequent players or about her own payoffs. That's why the types' beliefs are well defined irrespective of the players' actions, so optimal actions can be computed by backward induction, and forward-induction reasoning is irrelevant.

We will now describe this example formally, with deterministic game-positions and deterministic game-states. All the spaces we will consider are finite, so the σ-algebra of each of them will consist of all its subsets.[4]

The space P_0 of leaf game-positions is

$$\{-8, -1, 0, 1, 8\}^3$$

which describe the potential payoff-vectors to the players as envisaged by various types.

In the space P_1 the acting player is player 1, with the action set $A_1 = \{a_1, f_1\}$, so

$$P_1 = P_0^{A_1}$$

[4]In any case, we will not need these σ-algebras here because we consider only pure strategies.

Let

$$\tilde{P}_1 = P_0 \cup P_1$$

Similarly, P_2 is defined by

$$P_2 = \tilde{P}_1^{A_2}$$

where $A_2 = \{a_2, f_2\}$;

$$\tilde{P}_2 = P_2 \cup P_0;$$

and P_3 is defined by

$$P_3 = \tilde{P}_2^{A_3}$$

where $A_3 = \{a_3, f_3\}$.

Next, we describe the spaces of game-states.

$$\Omega_0 = P_0$$

is the space of leaf game-states.

$$\Omega_1 = \left\{(k_1, t_1) : k_1 \in K_1 = P_0^{A_1}, t_1 \in T_1 = \{0, 1\}\right\}$$

where

$$b_1 : T_1 \to \Delta(K_1)$$

is defined by:

- $b_1(t_1 = 1)$ assigns probability 1 to the function $k_1^1 \in K_1 = P_0^{A_1}$ defined by

$$k_1^1(a_1) = (8, 0, 1), \quad k_1^1(f_1) = (8, 8, 0) \qquad (k_1^1)$$

  (Recall that P_0 are payoff vectors for the three players. In fact, any distribution on K_1 would do as well, as long as it assigns probability 1 to functions $k_1 \in K_1$ with the property that $k_1(a_1)$ leads to payoff 1 to player 1, and $k_1(f_1)$ leads to payoff 0 to player 1. Player 1 – the granddaughter – does not form beliefs about the types of the other players, and she does not care about their payoffs.)

- $b_1(t_1 = 0)$ assigns probability 1 to the function $k_1^0 \in K_1 = P_0^{A_1}$ with the property that

$$k_1^0(a_1) = (8, 0, -1), \quad k_1^0(f_1) = (8, 8, 0) \qquad (k_1^0)$$

(A similar comment applies here, as well as to the definition of b_2 and b_3 in the sequel).

Next, define

$$\tilde{\Omega}_1 = \Omega_0 \cup \Omega_1$$

$$\Omega_2 = \left\{ (k_2, t_2) : k_2 \in K_2 = \tilde{\Omega}_1^{A_2}, t_2 \in T_2 = \{0,1\}^2 \right\}$$

where

$$b_2 : T_2 \to \Delta(K_2)$$

is defined by:

- $b_2(t_2 = 11)$ assigns probability $1-\varepsilon$ to the function $k_2^{11} \in K_2 = \tilde{\Omega}_1^{A_2}$ defined by

  $$k_2^{11}(a_2) = (0,1,1), \quad k_2^{11}(f_2) = \left(k_1^1, t_1 = 1\right) \tag{k_2^{11}}$$

  and probability ε to the function $\hat{k}_2^{11} \in K_2 = \tilde{\Omega}_1^{A_2}$ defined by

  $$\hat{k}_2^{11}(a_2) = (0,1,1), \quad \hat{k}_2^{11}(f_2) = \left(k_1^1, t_1 = 0\right) \tag{$\hat{k}_2^{11}$}$$

  (Thus, the type $t_2 = 11$ is certain what payoff vector would result if she takes the action a_2; she is also certain that taking the action f_2 would lead to a game-state in which player 1 is called to play. In the latter case, she is certain that the game-position in the game-state would be k_1^1, but she is uncertain about the type t_1 there.) Similarly,

- $b_2(t_2 = 10)$ assigns probability $1-\varepsilon$ to the function $k_2^{10} \in K_2 = \tilde{\Omega}_1^{A_2}$ defined by

  $$k_2^{10}(a_2) = (0,1,1), \quad k_2^{10}(f_2) = \left(k_1^0, t_1 = 0\right) \tag{k_2^{10}}$$

  and probability ε to the function $\hat{k}_2^{10} \in K_2 = \tilde{\Omega}_1^{A_2}$ defined by

  $$\hat{k}_2^{10}(a_2) = (0,1,1), \quad \hat{k}_2^{10}(f_2) = \left(k_1^0, t_1 = 1\right) \tag{$\hat{k}_2^{10}$}$$

- $b_2(t_2 = 01)$ assigns probability $1-\varepsilon$ to the function $k_2^{01} \in K_2 = \tilde{\Omega}_1^{A_2}$ defined by

  $$\begin{gathered} k_2^{01}(a_2) = (0,-1,1), \\ k_2^{01}(f_2) = \left(\begin{array}{c} (k_1 : A_1 \to P_0 : k_1(a_1) = (8,0,1), k_1(f_1) = (8,-8,0)), \\ t_1 = 1 \end{array} \right) \end{gathered} \tag{k_2^{01}}$$

  and probability ε to the function $\hat{k}_2^{01} \in K_2 = \tilde{\Omega}_1^{A_2}$ defined by

  $$\begin{gathered} \hat{k}_2^{01}(a_2) = (1,-1,0), \\ \hat{k}_2^{01}(f_2) = \left(\begin{array}{c} (k_1 : A_1 \to P_0 : k_1(a_1) = (8,0,1), k_1(f_1) = (8,-8,0)), \\ t_1 = 0 \end{array} \right) \end{gathered} \tag{$\hat{k}_2^{01}$}$$

- $b_2(t_2 = 00)$ assigns probability $1-\varepsilon$ to the function $k_2^{00} \in K_2 = \tilde{\Omega}_1^{A_2}$ defined by

$$k_2^{00}(f_2) = \begin{pmatrix} k_2^{00}(a_2) = (0,-1,1), \\ (k_1 : A_1 \to P_0 : k_1(a_1) = (8,0,-1), k_1(f_1) = (8,-8,0)), \\ t_1 = 0 \end{pmatrix} \qquad (k_2^{00})$$

and probability ε to the function $\hat{k}_2^{00} \in K_2 = \tilde{\Omega}_1^{A_2}$ defined by

$$\hat{k}_2^{00}(f_2) = \begin{pmatrix} \hat{k}_2^{00}(a_2) = (0,-1,1), \\ (k_1 : A_1 \to P_0 : k_1(a_1) = (8,0,-1), k_1(f_1) = (8,-8,0)), \\ t_1 = 1 \end{pmatrix} \qquad (\hat{k}_2^{00})$$

The following diagram depicts the game states for players 1 and (the granddaughter and the daughter).

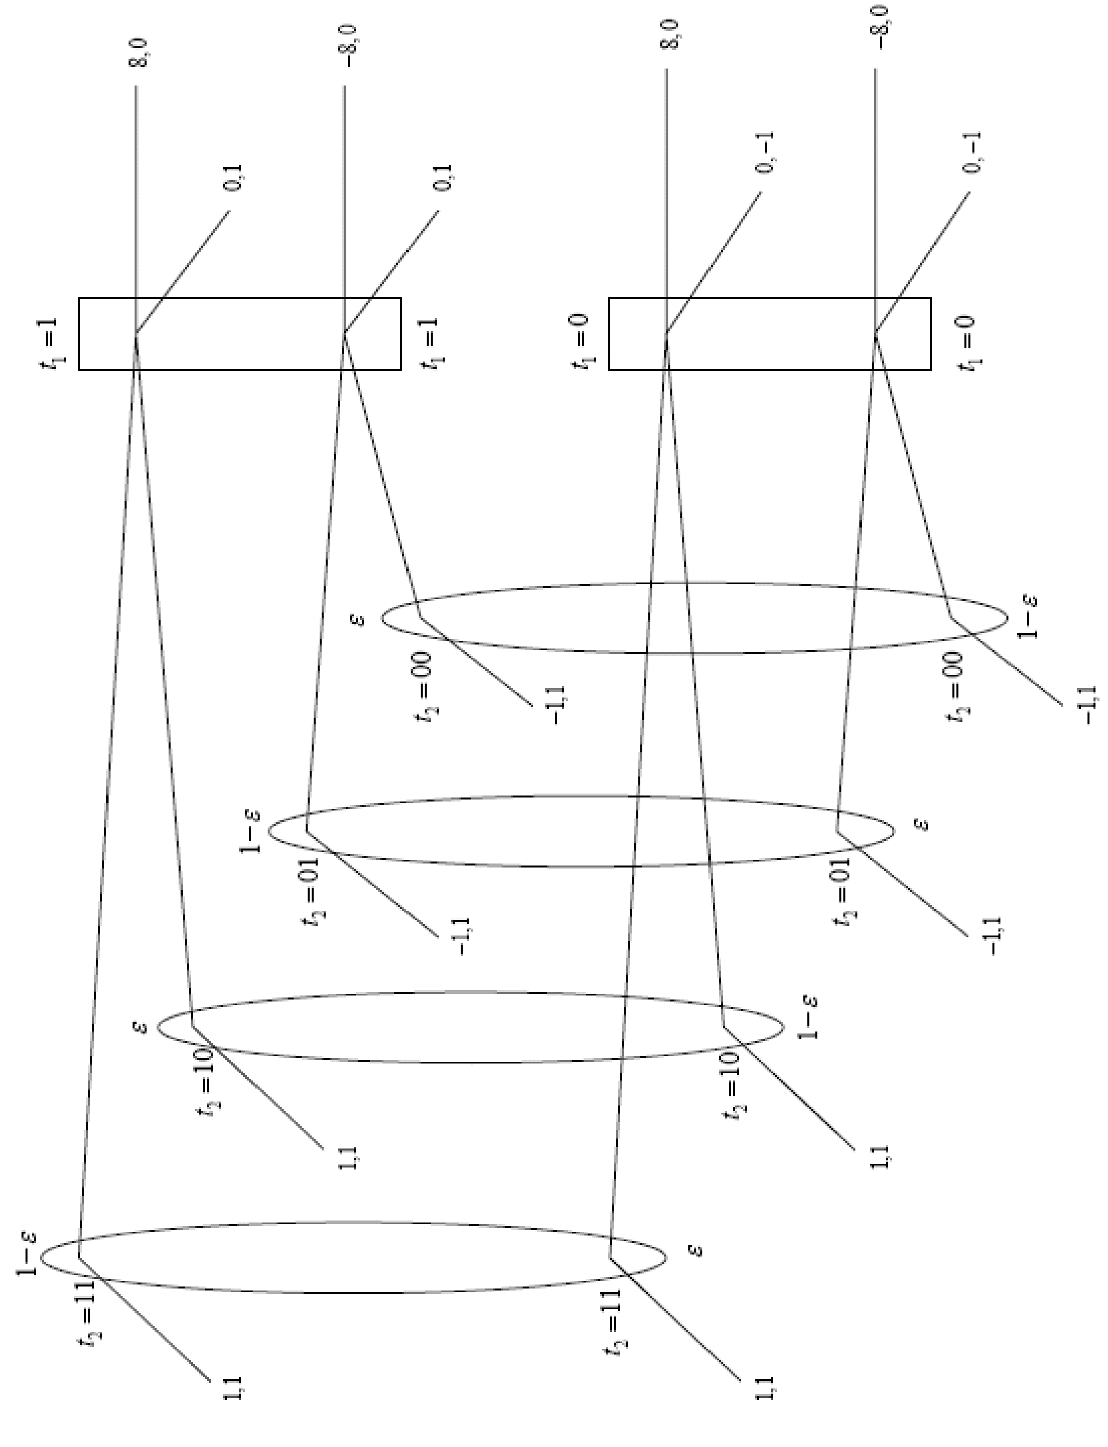
8,0
−8,0
8,0
−8,0
0,1
0,1
0,−1
0,−1
$t_1 = 1$
$t_1 = 1$
$t_1 = 0$
$t_1 = 0$
ε
1−ε
$t_2 = 00$
$t_2 = 00$
−1,1
−1,1
1−ε
ε
$t_2 = 01$
$t_2 = 01$
−1,1
−1,1
ε
1−ε
$t_2 = 10$
$t_2 = 10$
1,1
1,1
1−ε
ε
$t_2 = 11$
$t_2 = 11$
1,1
1,1

Last, define

$$\tilde{\Omega}_2 = \Omega_0 \cup \Omega_2$$

and

$$\Omega_3 = \left\{ (k_3, t_3) : k_3 \in K_3 = \tilde{\Omega}_2^{A_3}, t_3 \in T_3 = \{0,1\}^3 \right\}$$

where

$$b_3 : T_3 \to \Delta(K_3)$$

is defined by:

- $b_3(t_3 = 111)$ assigns probability $(1-\varepsilon)^2$ to the function $k_3^{111} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

  $$k_3^{111}(a_3) = (1,1,1), \quad k_3^{111}(f_2) = \left(k_2^{11}, t_2 = 11\right) \tag{k_3^{111}}$$

  probability $(1-\varepsilon)\varepsilon$ to the function $\hat{k}_3^{111} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

  $$\hat{k}_3^{111}(a_3) = (1,1,1), \quad \hat{k}_3^{111}(f_2) = \left(k_2^{11}, t_2 = 10\right) \tag{$\hat{k}_3^{111}$}$$

  probability $\varepsilon(1-\varepsilon)$ to the function $\check{k}_3^{111} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

  $$\check{k}_3^{111}(a_3) = (1,1,1), \quad \check{k}_3^{111}(f_2) = \left(k_2^{11}, t_2 = 01\right) \tag{$\check{k}_3^{111}$}$$

  and probability ε^2 to the function $\dot{k}_3^{111} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

  $$\dot{k}_3^{111}(a_3) = (1,1,1), \quad \dot{k}_3^{111}(f_2) = \left(k_2^{11}, t_2 = 00\right) \tag{$\dot{k}_3^{111}$}$$

- $b_3(t_3 = 110)$ assigns probability $(1-\varepsilon)^2$ to the function $k_3^{110} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

  $$k_3^{110}(a_3) = (1,1,-1), \quad k_3^{110}(f_2) = \left(k_2^{10}, t_2 = 10\right) \tag{k_3^{110}}$$

  probability $(1-\varepsilon)\varepsilon$ to the function $\hat{k}_3^{110} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

  $$\hat{k}_3^{110}(a_3) = (1,1,-1), \quad \hat{k}_3^{110}(f_2) = \left(k_2^{10}, t_2 = 11\right) \tag{$\hat{k}_3^{110}$}$$

  probability $\varepsilon(1-\varepsilon)$ to the function $\check{k}_3^{110} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

  $$\check{k}_3^{110}(a_3) = (1,1,-1), \quad \check{k}_3^{110}(f_2) = \left(k_2^{10}, t_2 = 00\right) \tag{$\check{k}_3^{110}$}$$

  and probability ε^2 to the function $\dot{k}_3^{110} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

  $$\dot{k}_3^{110}(a_3) = (1,1,-1), \quad \dot{k}_3^{110}(f_2) = \left(k_2^{10}, t_2 = 01\right) \tag{$\dot{k}_3^{110}$}$$

- $b_3(t_3 = 101)$ assigns probability $(1-\varepsilon)^2$ to the function $k_3^{101} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

$$k_3^{101}(a_3) = (1,-1,1), \quad k_3^{101}(f_2) = \left(k_2^{01}, t_2 = 01\right) \qquad (k_3^{101})$$

probability $(1-\varepsilon)\varepsilon$ to the function $\hat{k}_3^{101} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

$$\hat{k}_3^{101}(a_3) = (1,-1,1), \quad \hat{k}_3^{101}(f_2) = \left(k_2^{01}, t_2 = 00\right) \qquad (\hat{k}_3^{101})$$

probability $\varepsilon(1-\varepsilon)$ to the function $\check{k}_3^{101} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

$$\check{k}_3^{101}(a_3) = (1,-1,1), \quad \check{k}_3^{101}(f_2) = \left(k_2^{01}, t_2 = 11\right) \qquad (\check{k}_3^{101})$$

and probability ε^2 to the function $\dot{k}_3^{101} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

$$\dot{k}_3^{101}(a_3) = (1,-1,1), \quad \dot{k}_3^{101}(f_2) = \left(k_2^{01}, t_2 = 10\right) \qquad (\dot{k}_3^{101})$$

- $b_3(t_3 = 100)$ assigns probability $(1-\varepsilon)^2$ to the function $k_3^{100} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

$$k_3^{100}(a_3) = (1,-1,-1), \quad k_3^{100}(f_2) = \left(k_2^{00}, t_2 = 00\right) \qquad (k_3^{100})$$

probability $(1-\varepsilon)\varepsilon$ to the function $\hat{k}_3^{100} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

$$\hat{k}_3^{100}(a_3) = (1,-1,-1), \quad \hat{k}_3^{100}(f_2) = \left(k_2^{00}, t_2 = 01\right) \qquad (\hat{k}_3^{100})$$

probability $\varepsilon(1-\varepsilon)$ to the function $\check{k}_3^{100} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

$$\check{k}_3^{100}(a_3) = (1,-1,-1), \quad \check{k}_3^{100}(f_2) = \left(k_2^{00}, t_2 = 10\right) \qquad (\check{k}_3^{100})$$

and probability ε^2 to the function $\dot{k}_3^{100} \in K_3 = \tilde{\Omega}_2^{A_3}$ defined by

$$\dot{k}_3^{100}(a_3) = (1,-1,-1), \quad \dot{k}_3^{100}(f_2) = \left(k_2^{00}, t_2 = 11\right) \qquad (\dot{k}_3^{100})$$

What are the beliefs $b_3(t_3 = 011)$, $b_3(t_3 = 010)$, $b_3(t_3 = 001)$, $b_3(t_3 = 000)$? The same as the beliefs $b_3(t_3 = 011)$, $b_3(t_3 = 010)$, $b_3(t_3 = 001)$, $b_3(t_3 = 000)$, respectively, with the only difference that the payoffs to player 3 are multiplied by -1 in all the 4 game-states which are in the support of the type's belief[5], as well in all the game-states that emerge from these 4 game-states. We will skip the explicit description.

[5]With disjoint supports for different types!

4.2 Cheap Talk

We will now describe a simple cheap talk game, and one of its perfect Bayesian equilibria. Then we will describe the game states in the example. The example, though simple, should make clear that any perfect Bayesian equilibrium can be represented by a dynamic equilibrium with game-states.

Three major aspects make this example different from the example in the previous sub-section. First, the game starts with a chance move, defining the types of all the players right from the start. Thus, even the player who moves second has a well defined type and beliefs when the first player is called upon to play[6]. Second, as a result of the chance move, the types of the players originate from a common prior, and the second player, when called upon to play, updates her initial belief using that prior. Third, this update depends crucially also on the equilibrium behavior of the first player – under a different equilibrium the update would be different. Thus, game-states are part of the description of a perfect Bayesian equilibrium, and they do not exist independently of it or prior to it.

There are two states of nature – g ("good") and b ("bad") with equal likelihood. There are two players, $i = 1, 2$. They simultaneously get independent noisy signals $s_i \in [0, 1]$ about the state of nature. When the state of nature is g, then s_i is distributed so that $s_i = x$ with density $f_g(x) = 2x$. When the state of nature is b, then s_i is distributed so that $s_i = x$ with density $f_b(x) = 2(1 - x)$. Thus, $s_i = x$ is the posterior probability that player i assigns to the state of nature g.

Player 1 moves first. She can report either h ("high") or ℓ ("low"). Subsequently, player 2 can choose either r ("refrain") or a ("advance"). If player 2 chooses r, both players get a payoff of 0 irrespective of the state. If player 2 chooses a, then both players get a payoff -1 in the bad state b; in the good state g player 1 gets a payoff 2, and player 2 gets a payoff 1. Both players are expected payoff maximizers. In particular, player 2 would like to choose a if and only if she assigns posterior probability at least $\frac{1}{2}$ that the state of nature is g.

The game has several perfect Bayesian equilibria. One of them has the following structure. Player 1 reports h when her signal s_1 is larger than t; she reports ℓ otherwise. With this in mind, player 2 realizes that:

(1) If the state of nature is g, then player 2 would get the signal $s_2 = y$ and hear the message ℓ with density

$$y \int_0^t f_g(x)\, dx = y \int_0^t 2x dx = yt^2$$

[6] Though the first player doesn't know what this type is.

(2) If the state of nature is b, then player 2 would get the signal $s_2 = y$ and hear the message ℓ with density

$$(1-y)\int_0^t f_b(x)\,dx = (1-y)\int_0^t 2(1-x)\,dx = (1-y)\,t\,(2-t)$$

Thus, given the message ℓ, the posterior probability that player 2 assigns to the state g is

$$\frac{yt^2}{yt^2 + (1-y)\,t\,(2-t)}$$

which is larger than $\frac{1}{2}$ when $y > 1 - \frac{t}{2}$. Therefore, if player 2 hears the message ℓ, she chooses a whenever

$$s_2 > 1 - \frac{t}{2}$$

(3) If the state of nature is g, then player 2 would get the signal $s_2 = y$ and hear the message h with density

$$y\int_t^1 f_g(x)\,dx = y\int_t^1 2x dx = y\left(1 - t^2\right)$$

(4) If the state of nature is b, then player 2 would get the signal $s_2 = y$ and hear the message h with density

$$(1-y)\int_t^1 f_b(x)\,dx = (1-y)\int_t^1 2(1-x)\,dx = (1-y)(1-t)^2$$

Thus, given the message h, the posterior probability that player 2 assigns to the state g is

$$\frac{y(1-t^2)}{y(1-t^2) + (1-y)(1-t)^2}$$

which is larger than $\frac{1}{2}$ when $y > \frac{1-t}{2}$. Therefore, if player 2 hears the message h, she chooses a whenever

$$s_2 > \frac{1-t}{2}$$

This means that if player 1 reports ℓ, then when the state of nature is g, player 2 will choose a with probability

$$\int_{1-\frac{t}{2}}^1 f_g(y)\,dy = \int_{1-\frac{t}{2}}^1 2y dy = t - \frac{1}{4}t^2$$

and when the state of nature is b, player 2 will choose a with probability

$$\int_{1-\frac{t}{2}}^1 f_b(y)\,dy = \int_{1-\frac{t}{2}}^1 2(1-y)\,dy = \frac{t^2}{4}$$

Thus, if player 1 received the signal $s_1 = x$, she assigns probability x to the state g, so reporting ℓ yields her the expected payoff

$$U_1(x, \ell) = 2x\left(t - \frac{1}{4}t^2\right) + (-1)(1-x)\frac{t^2}{4}$$

Alternatively, if player 1 reports h, then when the state of nature is g, player 2 will choose a with probability

$$\int_{\frac{1-t}{2}}^{1} f_g(y)\,dy = \int_{\frac{1-t}{2}}^{1} 2y dy = \frac{(t+1)(3-t)}{4}$$

and when the state of nature is b, player 2 will choose a with probability

$$\int_{\frac{1-t}{2}}^{1} f_b(y)\,dy = \int_{\frac{1-t}{2}}^{1} 2(1-y)\,dy = \frac{(t+1)^2}{4}$$

With the signal $s_1 = x$, reporting h thus yields player 1 the expected payoff

$$U_1(x, h) = 2x\frac{(t+1)(3-t)}{4} + (-1)(1-x)\frac{(t+1)^2}{4}$$

We conclude that with the signal $s_1 = x^*$ player 1 is indifferent between reporting h or ℓ if

$$U_1(x^*, h) = U_1(x^*, \ell)$$

i.e. if

$$2x^*\left(t - \frac{1}{4}t^2\right) + (-1)(1-x^*)\frac{t^2}{4} = 2x^*\frac{(t+1)(3-t)}{4} + (-1)(1-x^*)\frac{(t+1)^2}{4}$$

At equilibrium $t = x^*$, and substitution in the equation above yields

$$t = \frac{5}{4} - \frac{1}{4}\sqrt{17} = 0.219\,22 \qquad (*)$$

One can then verify that indeed for $x \in (t, 1]$ we have

$$U_1(x, h) > U_1(x, \ell)$$

and for $x \in [0, t)$ we have

$$U_1(x, \ell) > U_1(x, h)$$

We are now going to describe the game states in this example. The potential payoffs for the players are

$$\begin{aligned} L_1 &= \{-1, 0, 2\} \\ L_2 &= \{-1, 0, 1\} \end{aligned}$$

The space of leaf game-positions, which is also the basic space of game-states, is

$$\Omega_0 = L = L_1 \times L_2$$

Next, in Ω_1 only player 2 acts – her action set is $A_2 = \{r, a\}$. There are two deterministic functions in $K = \{k_g, k_b\}$ specifying the players' payoffs as a function of player 2's action in A_2 in the corresponding state of nature. These functions are defined by

$$\begin{aligned} k_g(a) &= (2,1) \quad k_g(r) = (0,0) \\ k_b(a) &= (-1,-1) \quad k_g(r) = (0,0) \end{aligned}$$

The space of types of player 2 is

$$T_2 = T_2^h \cup T_2^\ell$$

where

$$T_2^h = T_2^\ell = [0,1]$$

These correspond to types who heard the message h or ℓ, respectively, though these messages are only implicit in the definition of the space of game-states Ω_1.

What are the beliefs of the different types of player 2? Here is a crucial difference in comparison with the previous example: The beliefs of the types correspond to the posterior beliefs of the types *at the particular perfect Bayesian equilibrium we are analyzing.* For example, there exists also another perfect Bayesian equilibrium in this game, in which player 1 never uses the message ℓ. Under this equilibrium, the space of types T_2 would be the same as above, *but the beliefs of the types would be completely different.*

The beliefs of the different types of player 2

$$b_2 : T_2^h \cup T_2^\ell \to \Delta(K)$$

are defined as follows by the probability assigned to the payoff function k_g of the good state:

$$\begin{aligned} b_2\left(t_2^\ell\right)(k_g) &= \frac{t_2^\ell t^2}{t_2^\ell t^2 + \left(1 - t_2^\ell\right) t (2 - t)}, \quad t_2^\ell \in T_2^\ell \\ b_2\left(t_2^h\right)(k_g) &= \frac{t_2^h \left(1 - t^2\right)}{t_2^h \left(1 - t^2\right) + \left(1 - t_2^h\right)(1 - t)^2}, \quad t_2^h \in T_2^h \end{aligned}$$

where t is the threshold type specified in $(*)$ above. This completes the definition of the space of game-states Ω_1.

We are now ready to define the space Ω_2 of game-states in which player 1 acts. Her action set is $A_1 = \{\ell, h\}$.

$$\Omega_2 = K' \times T_1$$

where

$$T_1 = [0,1]$$

and

$$K' = \{k_{g,y}, k_{b,y}\}_{y \in [0,1]}$$

consists of deterministic functions defined as follows:

$$k_{g,y} : A_1 \to \Omega_1$$

is defined by

$$\begin{aligned} k_{g,y}(h) &= (k_g, y) \in K \times T_2^h \\ k_{g,y}(\ell) &= (k_g, y) \in K \times T_2^\ell \end{aligned}$$

and

$$k_{b,y} : A_1 \to \Omega_1$$

is defined by

$$\begin{aligned} k_{b,y}(h) &= (k_b, y) \in K \times T_2^h \\ k_{b,y}(\ell) &= (k_b, y) \in K \times T_2^\ell \end{aligned}$$

K' is thus isomorphic to $\{g, b\} \times [0,1]$. The beliefs

$$b_1 : T_1 \to \Delta(K')$$

are defined as follows: The marginal of $b_1(t_1)$ on $\{g, b\}$ assigns probability t_1 to g and $1 - t_1$ to b. Conditional on g, the belief of $b_1(t_1)$ on $[0,1]$ is distributed with density $f_g(y) = 2y$. Conditional on b, the belief of $b_1(t_1)$ on $[0,1]$ is distributed with density $f_b(y) = 2(1-y)$.

This completes the definition of Ω_2.

Last, we define the space $\Omega_3 = \{\omega_3\}$ which consists of a unique game-state. This state is where nature acts: The space of active players is empty, so $A = \{\emptyset\}$, and $\omega_3 = k$, where k is the (non-deterministic) function

$$k : A \to \Delta(\Omega_2)$$

defined as follows. Recall that $\Omega_2 = K' \times T_1$, and K' is isomorphic to $\{g, b\} \times [0,1]$. Then $k(\emptyset)$ is a probability measure over $\{g, b\} \times [0,1] \times T_1$, such that the marginal on $\{g, b\}$ assigns probability $\frac{1}{2}$ to each of g and b; given g, $k(\emptyset)$ is a product distribution on $[0,1] \times T_1$, with density f_g on each coordinate; and given b, $k(\emptyset)$ is a product distribution on $[0,1] \times T_1$, with density f_b on each coordinate.

Thus, the beliefs $b_1(t_1)$ of the types of player 1 in Ω_2 are the posteriors of the prior on Ω_2 defined by $k(\emptyset)$. As we have already mentioned, the beliefs $b_2(t_2)$ of the types of player 2 in Ω_1 are further posteriors given the actions of player 1 *at this particular equilibrium.*

5 Variations and Generalizations

5.1 Information Sets

The apparatus of types in game-states can capture the classical notion of *information sets* in extensive-form games with imperfect information, once we define the belief of a type to be a *possibility set* rather than a probability measure. To this effect, (4) should be replaced by

$$b_i : T_i \longrightarrow 2^{K \times T_{-i}} \tag{4'}$$

With this definition, the information set of player i in a game-state $\omega = \left(k, (t_j)_{j \in I_0}\right)$ is

$$h_i \left(k, (t_j)_{j \in I_0}\right) = \left\{ \left(k', t_i, t'_{-i}\right) \in \Omega : \left(k', t'_{-i}\right) \in b_i (t_i) \right\}$$

where $k' \in K$, $t_i \in T_i$ and $t'_{-i} \in T_{-i} = \prod_{j \in I_0 \setminus i} T_j$.

Note that in the classical definition of extensive-form games with imperfect information, only a single player moves in each node, and hence T_{-i} is empty. Our definition allows for a more general setting, in which this restriction is not imposed.

Given the re-formulation of game states with (4'), let us restrict attention to classical game-states (which do not have two distinct paths ending in the same game-state). Notice that even under this restriction, a game-state can correspond to a node in an extensive-form game with *imperfect recall*. This is the case when two distinct paths that start in the same game-state ω end, respectively, in game-states ω', ω'' which belong to the same information set of some player i, even though the information sets of player i in the distinct paths did not intersect.

Nevertheless, a second kind of imperfect recall – absentmindedness – is excluded: An information set cannot contain two distinct game-states in a path of game-states. This is because the spaces of game-states are defined an inductive fashion, and if ω precedes ω' in a path then ω' has been defined earlier than ω in the inductive definition. That's why the definition of types in ω' cannot refer to ω, and in particular ω cannot belong to an information set containing ω'.

In order to allow for information sets which contain distinct game-states along a path, a "concurrent" rather then an "inductive" definition of game-states is needed. Such a generalization will be introduced in subsection 5.4 below.

5.2 Unawareness of Players and Actions

In the inductive definition of game-state spaces of section 3, we have defined one such space Ω at a time. In this way we made sure that the set of players I_0 and the action set A_i of each of these players $i \in I_0$ were the same across all the game-states of Ω. As a result, even if a type is uncertain regarding the game-state, she is nevertheless certain which actions are available to her, which other players play concurrently with her, and what actions each of them can take.

To allow for unawareness of some of the available actions or of some of the concurrent players, we introduce now a generalization of the inductive definition of game-states. Instead of defining one space Ω at a time, we will allow for defining simultaneously several $\tilde{\Omega}$-based game-state spaces, which differ by the identity of the active players and by their action sets. In a game-state ω the beliefs of a type of a player i can be concentrated on another $\tilde{\Omega}$-based game-state space Ω', as long as in Ω'

- Player i is one of the players ("cogito ergo sum")
- The set of players in Ω' is contained in the set of players in ω
- The set of actions available to the players in Ω' is a subset of the set of actions available to them in ω.

Here is the formal inductive definition.

The space of leaf game-positions is a space of game-states.

Let $\tilde{\Omega}$ be a set measurable space of game-states.
Let $\{I_0^\gamma\}_{\gamma\in\Gamma}$, $I_0^\gamma \subseteq I$ be a set of subsets of players.
Let $\left\{(A_i^\gamma)_{i\in I_0^\gamma}\right\}_{\gamma\in\Gamma}$ be measurable spaces of actions for the players $i \in I_0^\gamma$, $\gamma \in \Gamma$. Denote by $A^\gamma = \prod_{i\in I_0} A_i^\gamma$ the space of joint actions of the players in I_0^γ.
Let K^γ be a space of measurable functions

$$k^\gamma : A^\gamma \to \Delta\left(\tilde{\Omega}\right) \tag{3"}$$

from joint actions A^γ to probability distributions over game-states in $\tilde{\Omega}$. Denote $K = \bigcup_{\gamma\in\Gamma} K^\gamma$.
Let T_i^γ be a measurable space of *types* for each player $i \in I_0^\gamma$. Denote $T_{-i}^\gamma = \prod_{j\in I_0^\gamma\setminus i} T_j^\gamma$. Let

$$b_i^\gamma : T_i^\gamma \to \Delta\left(K^\gamma \times T_{-i}^\gamma\right) \tag{4"}$$

be a measurable function that specifies a probabilistic belief of each type $t_i^\gamma \in T_i^\gamma$ on K^γ and the other players' types T_{-i}^γ. Denote $T_i = \bigcup_{\gamma\in\Gamma} T_i^\gamma$

Then

$$\Omega^{\gamma}=\left\{\begin{array}{c}\omega^{\gamma}=\left(k^{\gamma},\left(t_i^{\delta_i}\right)_{i\in I_0^{\gamma}}\right)\in K^{\gamma}\times\prod_{i\in I_0^{\gamma}}T_i:\\ t_i^{\delta_i}\in T_i^{\delta_i}\implies i\in I_0^{\delta_i}\subseteq I_0^{\gamma},\quad A^{\delta_i}\subseteq A^{\gamma}\end{array}\right\}$$

is a space of $\tilde{\Omega}$-based game-states.

Thus, in a state $\omega^{\gamma}\in\Omega^{\gamma}$ the actual set of players is I_0^{γ} and their actual set of joint actions is A^{γ} (and so k^{γ} determines the distribution of game-states in $\tilde{\Omega}$ that will result from the actions of the players in the set I_0^{γ}). However, a player $i\in I_0^{\gamma}$ may believe that the set of players is actually $I_0^{\delta_i}\subseteq I_0^{\gamma}$ (as long as this player set $I_0^{\delta_i}$ contains i herself), and that the set of action profiles is $A^{\delta_i}\subseteq A^{\gamma}$. In such a case, player i is unaware of the players $j\in I_0^{\gamma}\setminus I_0^{\delta_i}$, and unaware of action profiles in $A^{\gamma}\setminus A^{\delta_i}$.

Consider the following example. An insurance company is obliged to sell an insurance policy against some major risk M at a regulated price, yielding a low expected payoff $L>0$ to the insurance company and a high expected utility h to the consumer. The insurance company sells also a combined insurance policy against the major risk M and a minor risk m at an unregulated price, with a high expected payoff H to the insurance company and a low expected utility ℓ to the consumer.

A person is initially aware of neither of the insurance options, and believes that her action set (vis-a-vis the risks) is empty. If she were aware of the insurance policies, she would prefer the regulated policy over the combined policy, and the combined policy over no-insurance: $h>\ell>0$.

The insurance company knows this, and can make the person aware of insurance options by posting an ad, at a negligible cost. The ad can refer to either of the insurance policies or to both (and the insurance company can also refrain from posting the ad altogether).

It should be clear how to describe the example with the above apparatus. At a dynamic equilibrium, the optimal action for the insurance company is to post an ad referring only to the combined policy. As a result, the person will believe that she has only two available actions – to buy the combined policy or to buy no insurance at all. This is the case even though she actually does have also a third available action – to buy the regulated policy. Given her belief, she will buy the combined policy, yielding the highest profit to the insurance company.

The following diagram depicts the game states in this example. The person can have one out of 4 possible types – $t_{\emptyset}$, t_M, t_{Mm}, t_{Mm+m} – depending on the content of the ad. The type's name reflects the set of insurance policies of which she is aware, and each of these types considers as possible a unique game-state, in which it is feasible for her to buy only these insurance policies By definition, each type takes an action, and the optimal action for each type is depicted by an arrow.

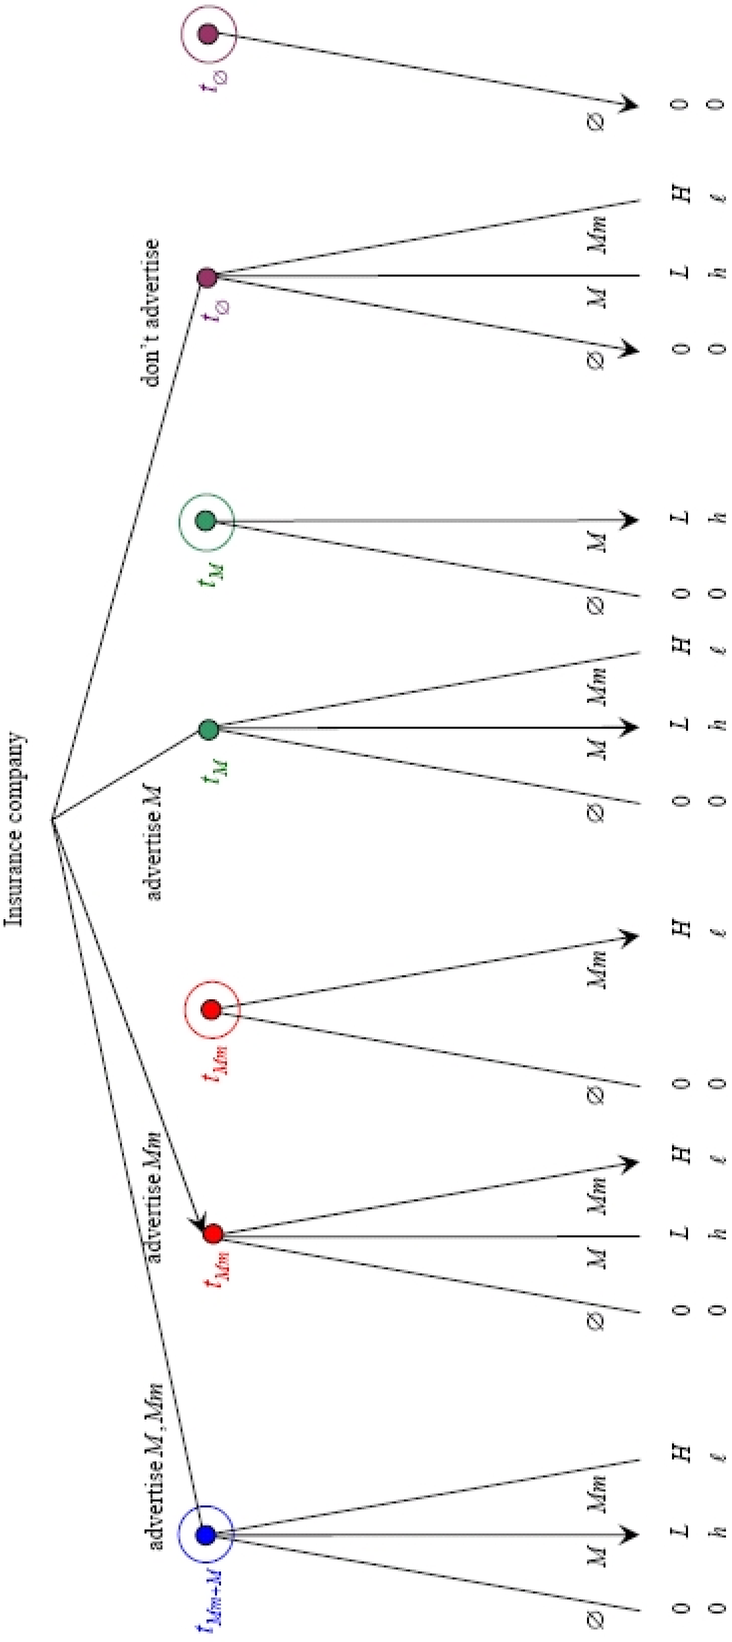
Insurance company
don't advertise
advertise M
advertise Mm
advertise M, Mm
$t_\varnothing$
t_M
t_{Mm}
t_{Mm+M}

5.3 Unawareness of aspects of reality

Consider the following scene from the film "The Princess Bride" by Rob Reiner (1987), based on a novel by William Goldman.[7] Westley, in the guise of the Dread Pirate Roberts, confronts his foe-for-the-moment, the Sicilian, Vizzini. Westley challenges him to a Battle of Wits. Two glasses are placed on the table, each containing wine and one purportedly containing poison. The challenge, simply, is to select the glass that does not lead to immediate death.

> Roberts: All right: where is the poison? The battle of wits has begun. It ends when you decide and we both drink, and find out who is right and who is dead.
>
> Vizzini: But it's so simple. All I have to do is divine from what I know of you. Are you the sort of man who would put the poison into his own goblet, or his enemy's? Now, a clever man would put the poison into his own goblet, because he would know that only a great fool would reach for what he was given. I'm not a great fool, so I can clearly not choose the wine in front of you. But you must have known I was not a great fool; you would have counted on it, so I can clearly not choose the wine in front of me.
>
> Roberts: You've made your decision then?
>
> Vizzini: Not remotely. Because iocane comes from Australia, as everyone knows. And Australia is entirely peopled with criminals. And criminals are used to having people not trust them, as you are not trusted by me. So I can clearly not choose the wine in front of you.
>
> ...
>
> Roberts: You're trying to trick me into giving away something – it won't work –
>
> Vizzini: (triumphant) It has worked – you've given everything away – I know where the poison is.
>
> Roberts: (fool's courage) Then make your choice.

[7] The following introduction to the scene was compiled by Mike Shor, http://www.gametheory.net/popular/reviews/PrincessBride.html
A video clip of the scene is posted at
http://www.gametheory.net/media/Princess.wmv
The full script of the film is posted at
http://www.godamongdirectors.com/scripts/princess.shtml

Vizzini: I will. And I choose – (And suddenly he stops, points at something behind Roberts) – what in the world can that be?

Roberts: (turning around, looking) What? Where? I don't see anything.

Vizzini: (busily switching the goblets while Roberts has his head turned.) Oh, well, I-I could have sworn I saw something. No matter.

(Roberts to face him again. Vizzini starts to laugh.)

Roberts: What's so funny?

Vizzini: I'll tell you in a minute. First, let's drink – me from my glass, and you from yours.

(And he picks up his goblet. Roberts picks up the one in front of him. As they both start to drink, Vizzini hesitates a moment. Then, allowing Roberts to drink first, he swallows his wine.)

Roberts: You guessed wrong.

Vizzini: (roaring with laughter) You only think I guessed wrong – (louder now) – that's what's so funny! I switched glasses when your back was turned. You fool. ... You fell victim to one of the classic blunders. The most famous is "Never get involved in a land war in Asia." But only slightly less well known is this: "Never go in against a Sicilian when death is on the line."
(He laughs and roars and cackles and whoops and is in all ways quite cheery until he falls over dead.)

Roberts: (stepping past the corpse, taking the blindfold and bindings off Buttercup [the girl over whom this battle was staged in the first place] She notices Vizzini lying dead. Roberts pulls her to her feet.)

...
Buttercup: To think – all that time it was your cup that was poisoned.

Roberts: They were both poisoned. I spent the last few years building up an immunity to iocane powder.

There are probably many ways to model the strategic interaction in this scene. We will be making the following assumptions:

- Clearly, Vizzini has no clue which glass is poisoned, so we will assume he initially ascribes equal probabilities to the two options. He assigns probability zero to the case which turned to be true, by which both glasses were poisoned.
- Vizzini is unaware of the possibility that somebody (and Roberts in particular) could be immune to the poison. This unawareness is not about an available action, but rather about an aspect of reality.
- When Roberts turned around, he could conceivably be unaware that Vizzini is able to switch the glasses. This unawareness is thus about an action available to Vizzini. In the analysis that follows, we will assume, for the sake of simplicity, that Roberts is indeed unaware that Vizzini could have switched the glasses. (In an appendix, we detail the analysis assuming that Vizzini is initially unsure whether or not Roberts is aware of Vizzini's ability to switch the glasses, and assigns equal probabilities to the two possibilities – awareness and unawareness.)
- If Roberts drinks from his glass, then Vizzini has no option but to follow suit (we rule out the physical option that Vizzini has to simply refrain from drinking – we assume this is not an action available to him). If, alternatively, Roberts hesitates, Vizzini doesn't drink either.

We assume further that the payoffs are 0 to both players if none of them drinks, -2 to a player who dies (this will be a "negative-enough" payoff for our purpose here), and 1 to a player who stays alive while the other one dies. So $L_i = \{-2, 0, 1\}$, $i = 1, 2$, and we define $L = L_1 \times L_2$.

The last player to act is Roberts – we designate him as player 1. He essentially decides whether they both drink – dd – or both refrain – rr – from drinking, so

$$A_1 = \{dd, rr\}$$

If they refrain from drinking, they both get 0 irrespective of the content of the glasses. If they both drink, their payoffs depend on the content of the glasses, on whether the glasses have been switched, and on whether Roberts is immune to the poison.

Thus, in

$$K = L^{A_1}$$

there are only 4 relevant payoff functions – corresponding to the identity of the surviving players if they both drink:

$$k_{+-}(dd) = (1, -2), \quad k_{+-}(rr) = (0, 0) \tag{k_{+-}}$$

$$k_{-+}(dd) = (-2, 1), \quad k_{-+}(rr) = (0, 0) \tag{k_{-+}}$$

$$k_{--}(dd) = (-2, -2), \quad k_{--}(rr) = (0, 0) \tag{k_{--}}$$

$$k_{++}(dd) = (0, 0), \quad k_{++}(rr) = (0, 0) \tag{k_{++}}$$

Roberts is aware of the fact that he could be immune to the poison (since in fact he is). Roberts may be either aware or unaware of the fact that Vizzini could have switched the glasses. In the former case, we assume that Roberts assigns equal probabilities to the two options – that the glasses have been switched or not switched.

However, Vizzini is unaware of the immunity issue, *so in Vizzini's mind* the types of Roberts have nothing to do with immunity.

Thus, altogether we should consider *2 different sets of types* of Roberts. In the following specification of types in $\ddot{T}_1$, the subscript represents the content of the glasses in their original order: "1" represents poison, "0" represents no-poison. The left digit in the subscript refers to the glass initially in front of Roberts, the right digit refers to the glass in front of Vizzini. The superscript has the letter "i" if Roberts is $i$mmune to the poison, and the letter "v" if Roberts is $v$ulnerable to the poison. However, in $\mathring{T}_1$ the superscript is missing (describing types of Roberts *in the mind of Vizzini*)

$$\begin{aligned}\ddot{T}_1 &= \{t_{10}^v, t_{01}^v, t_{11}^v, t_{10}^i, t_{01}^i, t_{11}^i\} \\ \mathring{T}_1 &= \{t_{10}, t_{01}, t_{11}\}\end{aligned}$$

The beliefs of the different types are listed in the following tables. Each row in a table corresponds to a type, and each column corresponds to a function in K. The entries specify the probabilities that the corresponding types ascribe to the corresponding functions in K.

$\ddot{T}_1$	k_{+-}	k_{-+}	k_{--}	k_{++}
t_{10}^v	0	1	0	0
t_{01}^v	1	0	0	0
t_{11}^v	0	0	1	0
t_{10}^i	0	0	0	1
t_{01}^i	1	0	0	0
t_{11}^i	1	0	0	0

$\mathring{T}_1$	k_{+-}	k_{-+}	k_{--}	k_{++}
t_{10}	1	0	0	0
t_{01}	0	1	0	0
t_{11}	0	0	1	0

Correspondingly, there are two spaces –

$$\begin{aligned}\ddot{\Omega} &= K \times \ddot{T}_1 \\ \mathring{\Omega} &= K \times \mathring{T}_1\end{aligned}$$

which form a lattice – $\mathring{\Omega}$ is a quotient space of $\ddot{\Omega}$. This lattice of spaces constitutes an *unawareness structure* – a generalized kind of a type space introduced by Heifetz, Meier and Schipper (2005, forthcoming) for modeling *mutual* awareness and unawareness.

In this particular example the lattice of spaces is very simple, because only Roberts acts and has types in this unawareness structure. If there had been several players acting simultaneously, the lattice of spaces would have been more involved. Anyway, this example is meant to demonstrate that when *unawareness about an aspect of reality* is involved, game-states will constitute an *unawareness structure* rather than a classical type space.

Notice that if it were only for Roberts, we wouldn't need $\mathring{\Omega}$ in the lattice. However, we now turn to describe Vizzini's state of mind when Roberts has just turned around, and Vizzini has to decide whether or not to switch the glasses. In Vizzini's mind, each of these actions would lead to a state of affairs described by a game-state in $\mathring{\Omega}$. However, game-states in $\mathring{\Omega}$ do not represent actual game situations – they do not capture all the aspects relevant to the strategic interaction. The truth of the matter is that Vizzini's action leads to a game-state in $\ddot{\Omega}$. That's why we need the entire unawareness structure consisting of the two spaces $\ddot{\Omega}$ and $\mathring{\Omega}$.

In Vizzini's mind there are 3 functions that could potentially determine the game-state as a result of Vizzini's action – Not to switch or to Switch

$$A'_2 = \{N, S\}$$

These functions

$$\check{K}' = \{k_{10}, k_{01}, k_{11}\}$$

are from A'_2 to $\mathring{\Omega}$: as a function of Vizzini's choice – N or S – they determine the payoff function in K and Roberts' type in $\mathring{T}_1$, as follows:

$$\begin{aligned} k_{10}(N) &= (k_{+-}, t_{10}), \quad k_{10}(S) = (k_{-+}, t_{10}) \\ k_{01}(N) &= (k_{-+}, t_{01}), \quad k_{01}(S) = (k_{+-}, t_{01}) \\ k_{11}(N) &= (k_{--}, t_{11}), \quad k_{11}(S) = (k_{--}, t_{11}) \end{aligned}$$

We now have to describe Vizzini's types. In fact, Vizzini has a single type –

$$T'_2 = \{t_2\}$$

with the belief

$$b_2(t_2)(k_{10}) = b_2(t_2)(k_{01}) = \frac{1}{2}$$

Thus, Vizzini is uncertain whether his glass or Roberts' glass is poisoned, and ascribes equal weights to both possibilities.

The truth of the matter is that reality is described in richer terms than in Vizzini's mind, by the functions in

$$\hat{K}' = \left\{k^v_{10}, k^v_{01}, k^v_{11}, k^i_{10}, k^i_{01}, k^i_{11}\right\}$$

from A'_2 to game-states in $\ddot{\Omega}$. They are defined as follows:

$$\begin{aligned}
k^v_{10}(N) &= (k_{-+}, t^v_{10}), \quad k^v_{10}(S) = (k_{+-}, t^v_{10}) \\
k^v_{01}(N) &= (k_{+-}, t^v_{01}), \quad k^v_{01}(S) = (k_{-+}, t^v_{01}) \\
k^v_{11}(N) &= (k_{--}, t^v_{11}), \quad k^v_{11}(S) = (k_{+-}, t^v_{10}) \\
k^i_{10}(N) &= \left(k_{++}, t^i_{10}\right), \quad k^i_{10}(S) = \left(k_{+-}, t^i_{10}\right) \\
k^i_{01}(N) &= \left(k_{+-}, t^i_{01}\right), \quad k^i_{01}(S) = \left(k_{++}, t^i_{01}\right) \\
k^i_{11}(N) &= \left(k_{+-}, t^i_{11}\right), \quad k^i_{11}(S) = \left(k_{+-}, t^i_{11}\right) \qquad (**)
\end{aligned}$$

Altogether, we now have an unawareness structure consisting of a lattice of two spaces –

$$\begin{aligned}
\hat{\Omega}' &= \hat{K}' \times T'_2 \\
\check{\Omega}' &= \check{K}' \times T'_2
\end{aligned}$$

where the latter is a quotient space of the former. The situation is described by one of the game-states in $\hat{\Omega}'$. However, in Vizzini's mind the situation is described by a game-state in $\check{\Omega}'$ (and with his type t_2 Vizzini ascribes equal probabilities to two of the game-states in $\check{\Omega}'$).

We can now compute, by backward induction, the optimal action for each of the types. The optimal actions for the types of Roberts in $\ddot{T}_1, \mathring{T}_1$ are easy to compute. In particular, let us specify first the optimal action $\sigma_1(\cdot)$ of the following types in $\mathring{T}_1$:

$$\begin{aligned}
\sigma_1(t_{10}) &= rr \\
\sigma_1(t_{01}) &= dd
\end{aligned}$$

Of these two types only t_{01} (who is certain that only Vizzini's glass is poisoned) chooses to drink.

This means that if we fold the game-states in $\mathring{\Omega}$ with this optimal strategy of Roberts, then

$$\begin{aligned}
&(k_{-+}, t_{10}) \text{ is replaced by the payoff vector } (0, 0) && \text{(e)} \\
&(k_{-+}, t_{10}) \text{ is replaced by the payoff vector } (0, 0) && \text{(f)} \\
&(k_{+-}, t_{01}) \text{ is replaced by the payoff vector } (1, -2) && \text{(g)} \\
&(k_{-+}, t_{01}) \text{ is replaced by the payoff vector } (-2, 1) && \text{(h)}
\end{aligned}$$

With his type t_2', Vizzini believes that choosing Not to switch the glasses will lead, with equal probabilities, to one of the two game-states (e),(g), and hence to a negative expected payoff:

$$\frac{1}{2} \times 0 + \frac{1}{2} \times (-2) = -1$$

while, according to this belief, choosing to Switch the glasses will lead, with equal probabilities, to one of the game-states (f),(h), and hence to a positive expected payoff to himself:

$$\frac{1}{2} \times 0 + \frac{1}{2} \times 1 = \frac{1}{2}$$

So the optimal action of (the unique type t_2 of) Vizzini is S.

However, the truth of the matter is that both glasses are poisoned and Roberts is immune to the poison – the true game-state is (k_{11}^i, t_2'). From $(**)$ it follows that when Vizzini chooses S (and in fact even if he were to choose N) the resulting payoff function is k_{+-}, and the type of Roberts is t_{11}^i. From the last line in the table defining the beliefs of the types in $\ddot{T}_1$, we see that Roberts is certain that the payoff function is k_{+-}. From the definition of k_{+-} we see that the optimal action for Roberts is dd.

The game states in this analysis are depicted in the following diagram.

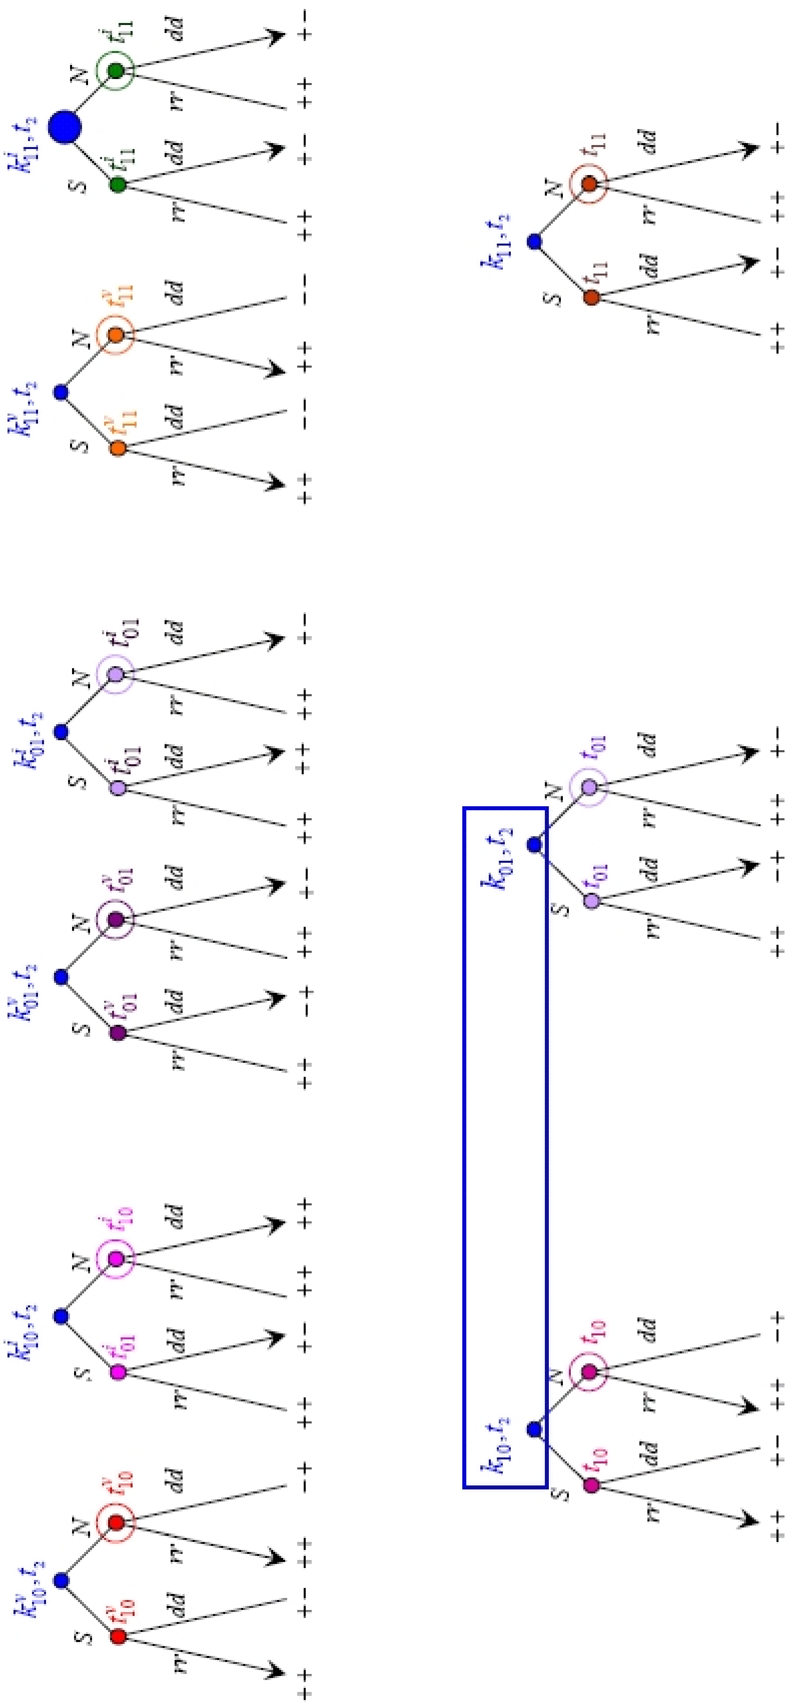

5.4 Generalized game-states and absentminded types

In section 5.1 we have already seen how one kind of imperfect recall can be captured by game-states – when a type of a player considers as possible (or assigns a positive probability to) two game-states that end distinct paths, while along each path the player excluded (or assigned zero probability to) game-states in the other path.

However, as explained there, a player cannot be absentminded – she cannot consider as possible (or assign a positive probability to) two distinct game-states along the same path. This is because the spaces of game-states are defined inductively, and if ω precedes ω' in a path then ω' has been defined earlier than ω in the inductive definition. That's why the definition of types in ω' cannot refer to ω. To allow for such types, we need a "concurrent" rather then an "inductive" definition of game-states. Such a definition of generalized game-states is presented here.

Let L be the measurable space of leaf game-positions, as defined in section 2. Let I be a set of players, and A_i a measurable space of actions for each player $i \in I$. Denote by $A = \prod_{i \in I} A_i$ the space of joint actions of the players.

A space of generalized game-states is a tuple

$$\left\langle \Omega, \kappa, (b_i)_{i \in I} \right\rangle$$

in which:

1. The set of generalized game-states Ω is
$$\Omega = L \cup \left(K \times \prod_{i \in I} T_i \right)$$
where K and the type sets T_i, $i \in I$ are all measurable spaces;

2. κ is a measurable function
$$\kappa : K \to (\Delta(\Omega))^A \tag{3"'}$$
that associates to each $k \in K$ a *measurable* function from joint actions in A to probability measures on Ω; or, alternatively,
$$\kappa : K \to \Omega^A \tag{5"'}$$
associates to each $k \in K$ a *measurable* function from joint actions in A to Ω, in which case we get a space of *generalized* ***deterministic*** *game-states;*

3.
$$b_i : T_i \to \Delta(K \times T_{-i}) \tag{4"'}$$
is a measurable function that specifies a probabilistic belief of each type $t_i \in T_i$ on K and the other players' types $T_{-i} = \prod_{j \neq i} T_j$; and

4. The game-positions in all the generalized game-states in Ω are identified by the following inductive definition of the map Π:

a) Π is the identity map on L:

$$\Pi(\ell) = \ell, \quad \ell \in L$$

b) If $\tilde{\Omega} \subset \Omega$ is the subset of generalized game-states on which Π is already defined, and $\Omega' \subseteq \Omega \setminus \tilde{\Omega}$ is the set of generalized game-states $\left(k, (t_i)_{i \in I}\right)$ in which $\kappa(k)(a) \in \Delta\left(\tilde{\Omega}\right)$ for all $a \in A$ (or, if κ is defined by (5"'), $\kappa(k)(a) \in \tilde{\Omega}$), then for each such $\left(k, (t_i)_{i \in I}\right) \in \Omega'$, the game-position $\Pi\left(\left(k, (t_i)_{i \in I}\right)\right)$ is defined by[8]

$$\Pi\left(\left(k, (t_i)_{i \in I}\right)\right)(a) = \Pi(\kappa(k)(a)), \qquad a \in A$$

Two features distinguish this "concurrent" definition of generalized game-states from the "inductive" definition of game-states in section 3. In the inductive definition, $k \in K$ were functions from joint actions $a \in A$ to game-states (or probability measures over game-states) in a previously-defined space $\tilde{\Omega}$ of game-states. Here, in contrast, the range of these functions is the entire space Ω itself, defined inter alia using K. Thus, in order for the definition not to be circular, we first have to define K as an abstract set of objects, and then to associate to each $k \in K$ a function from joint actions $a \in A$ to game-states (or probability measures over game-states). This association is defined by the function κ.

Hence, $\Omega \setminus L$ is a particular kind of a type space, where K is the space of states of nature, and κ can be thought of as a valuation function which describes the "content" of each state of nature $k \in K$. Here, this "content" $\kappa(k)$ is simply a description of how the game would evolve as a function of the players actions $a \in A$, i.e. what generalized game-state the players would find themselves in as a function of their actions $a \in A$ (or according to which distribution this generalized game-state would be drawn as a function of their actions $a \in A$).

Another difference is that here all the players have action sets and types, while in the inductive definition of section 3, only some (possibly proper) subset of the players had types and actions in each round of the inductive definition. If in that definition we would ascribe also to each "non-active" player a singleton action set and a singleton type set

[8] If case $\kappa(k)(a)$ is a probability measure in $\Delta\left(\tilde{\Omega}\right)$, the probability measure $\Pi(\kappa(k)(a))$ on game-positions is defined by

$$\Pi(\kappa(k)(a))(E) = \kappa(k)(a)\left(\Pi^{-1}(E)\right)$$

for all measurable sets E of game-positions.

(with an arbitrary belief), and define also a function κ on K to be the identity, the class of game-states would become a sub-class of the class of generalized game-states.

Also, to signify that some action $a_i \in A_i$ cannot be taken at a generalized game-state, it should be associated with a very negative payoff. The action set A_i of player i is thus the union of the relevant actions in the different generalized game-states.

Condition 4 ensures that there are no infinite paths in a space Ω of generalized game-states. This is because the function Π maps paths of generalized game-states to paths of game-positions in a one-to-one fashion, and there are no infinite paths of game-positions. (For the spaces of game-states defined in section 3, the finiteness of paths followed as a consequence of the inductive nature of the definition.)

Recall the Piccione and Rubinstein (1997) example of the absent-minded driver. Late at night in the bar, the driver knows that he should take the second exit on the highway in order to get home with a payoff of 4; continuing will lead to a motel with a payoff of 1. But the first exit, which leads to a disastrous neighborhood and a payoff of 0, looks identical, and the driver knows that upon reaching any of the intersections he will not be able to tell them apart, nor would he remember whether he has already passed an intersection on his way.

This story can be represented by generalized game-states in several ways. In what follows, we present the "simplest" formulation.[9]

Here $L_1 = \{-\infty, 0, 1, 4\}$, $A_1 = \{\text{continue, exit}\}$, $K = \{\text{bar, intersection\_1, intersection\_2}\}$, $T_1 = \{t_{\text{bar}}, t_{\text{intersections}}\}$

$$
\begin{aligned}
\kappa\,(\text{bar})\,(\text{exit}) &= -\infty \\
\kappa\,(\text{bar})\,(\text{continue}) &= (\text{intersection\_1}, t_{\text{intersections}}) \\
\kappa\,(\text{intersection\_1})\,(\text{exit}) &= 0 \\
\kappa\,(\text{intersection\_1})\,(\text{continue}) &= (\text{intersection\_2}, t_{\text{intersections}}) \\
\kappa\,(\text{intersection\_2})\,(\text{exit}) &= 4 \\
\kappa\,(\text{intersection\_2})\,(\text{continue}) &= 1
\end{aligned}
$$

The beliefs of the types are

$$
\begin{aligned}
b_1\,(t_{\text{bar}})\,(\text{bar}) &= 1 \\
b_1\,(t_{\text{intersections}})\,(\text{intersection\_1}) &= \alpha, \quad b_1\,(t_{\text{intersections}})\,(\text{intersection\_2}) = 1-\alpha
\end{aligned}
$$

This formulation of highlights the following issues:

[9] A different formulation would result by representing the multi-self game suggested by Gilboa (1997) with generalized game-states.

1. A strategy is defined by the choice of each type, not by one coordinated choice of a player for the entire set of his types in game-states in which he may be playing. In our setting it is meaningless to ask what would t_{bar} like to do if he were to choose at the game-state (intersection_1, $t_{\text{intersections}}$) – at that game-state it is the type $t_{\text{intersections}}$ which chooses and acts. t_{bar} has no choice but to take as given the strategy of $t_{\text{intersections}}$ at (intersection_1, $t_{\text{intersections}}$) and (intersection_2, $t_{\text{intersections}}$).

2. The type $t_{\text{intersections}}$ knows that it is he himself who chooses both in (intersection_1, $t_{\text{intersections}}$) and (intersection_2, $t_{\text{intersections}}$) , which he cannot tell apart and hence behaves in exactly the same way in both.[10]

3. The strategy "continue with probability p" yields to $t_{\text{intersections}}$ an expected payoff of
$$\alpha\left(4p\left(1-p\right)+p^{2}\right)+\left(1-\alpha\right)\left(4\left(1-p\right)+p\right)$$
which is maximized at
$$p=\max\left\{\frac{7\alpha-3}{6\alpha},0\right\}$$
The definition of generalized game-states does not restrict a priori the probability α in the belief of the type $t_{\text{intersections}}$. As part of a solution concept, one could add, for example, a requirement that this probability is coherent with the probability p with which $t_{\text{intersections}}$ chooses to continue:
$$\alpha=\frac{1}{1+p}$$
With this extra restriction, we get
$$\begin{aligned} p &= \frac{4}{9} \\ \alpha &= \frac{9}{13} \end{aligned}$$
which is one of the "solutions" suggested by Piccione and Rubinstein (1997). With this approach, the type $t_{\text{intersections}}$ takes his own belief $b_i\left(t_{\text{intersections}}\right)$ as given and chooses p so as to maximize his perceived expected payoff; it "so happens" that the belief is actually coherent with the choice.

4. Suppose that, as in section 5.1 above, we were to replace the probabilistic type $t_{\text{intersections}}$ by an "information set" type $\bar{t}_{\text{intersections}}$defined by the possibility set

[10] Thus, these game-states are not compatible with the multi-self approach discussed in section 7 of Piccione and Rubinsten (1997), and advocated by Aumann, Hart and Perry (1997). As Lipman (1997) notes, Gilboa's (1997) suggestion for a game between two agents of the same player does account for this approach; as mentioned above, that game is expressible with a space of generalized game-states which is different, of course, than the one presented here.

{intersection_1, intersection_2}. This would reflect, in fact, the original formulation of the story of Piccione and Rubinstein. With this formulation, the probability α is not defined in advance, and would be a consequence of the choice of p by $t_{\text{intersections}}$. This type would then be saying to himself: "The smaller the probability p with which I decide to continue, the smaller the probability

$$1-\alpha=\frac{p}{1+p}$$

with which I would find myself deciding in the game-state (intersection_2, $t_{\text{intersections}}$). Thus, if I choose to continue with probability p, my expected payoff would be

$$\begin{aligned} & \alpha\left[4p\left(1-p\right)+p^{2}\right]+\left(1-\alpha\right)\left[4\left(1-p\right)+p\right] \\ = & \frac{1}{1+p}\left[4p\left(1-p\right)+p^{2}\right]+\frac{p}{1+p}\left[4\left(1-p\right)+p\right] \end{aligned}$$

which is maximized at

$$p^{*}=\sqrt{\frac{7}{3}}-1=0.527\,53$$

(Note: I do not know if this "solution" has already been discussed in the literature which addressed this example).

6 Appendix: The Princess Bride scene – a richer model

In this appendix we provide a richer model for the "Princess Bride" scene, assuming that Roberts could potentially be aware that Vizzini may have switched the glasses. Here are the explicit assumptions

- Clearly, Vizzini has no clue which glass is poisoned, so we will assume he initially ascribes equal probabilities to the two options. He assigns probability zero to the case which turned to be true, by which both glasses were poisoned.
- Vizzini is unaware of the possibility that somebody (and Roberts in particular) could be immune to the poison. This unawareness is not about an available action, but rather about an aspect of reality.
- When Roberts turned around, he could conceivably be unaware that Vizzini is able to switch the glasses. This unawareness is thus about an action available to Vizzini. Vizzini is initially unsure whether or not Roberts is aware of Vizzini's ability to switch the glasses, and assigns equal probabilities to the two possibilities – awareness and unawareness. (in contrast, in section 5.3 we assumed that Roberts is unaware that Vizzini could have switched the glasses.)

- If Roberts drinks from his glass, then Vizzini has no option but to follow suit (we rule out the physical option that Vizzini has to simply refrain from drinking – we assume this is not an action available to him). If, alternatively, Roberts hesitates, Vizzini doesn't drink either.

We assume further that the payoffs are 0 to both players if none of them drinks, -2 to a player who dies (this will be a "negative-enough" payoff for our purpose here), and 1 to a player who stays alive while the other one dies. So $L_i = \{-2, 0, 1\}$, $i = 1, 2$, and we define $L = L_1 \times L_2$.

The last player to act is Roberts – we designate him as player 1. He essentially decides whether they both drink – dd – or both refrain – rr – from drinking, so

$$A_1 = \{dd, rr\}$$

If they refrain from drinking, they both get 0 irrespective of the content of the glasses. If they both drink, their payoffs depend on the content of the glasses, on whether the glasses have been switched, and on whether Roberts is immune to the poison.

Thus, in

$$K = L^{A_1}$$

there are only 4 relevant payoff functions – corresponding to the identity of the surviving players if they both drink:

$$k_{+-}(dd) = (1, -2), \quad k_{+-}(rr) = (0, 0) \tag{k_{+-}}$$
$$k_{-+}(dd) = (-2, 1), \quad k_{-+}(rr) = (0, 0) \tag{k_{-+}}$$
$$k_{--}(dd) = (-2, -2), \quad k_{--}(rr) = (0, 0) \tag{k_{--}}$$
$$k_{++}(dd) = (0, 0), \quad k_{++}(rr) = (0, 0) \tag{k_{++}}$$

Roberts is aware of the fact that he could be immune to the poison (since in fact he is). Roberts may be either aware or unaware of the fact that Vizzini could have switched the glasses. In the former case, we assume that Roberts assigns equal probabilities to the two options – that the glasses have been switched or not switched.

However, Vizzini is unaware of the immunity issue, *so in Vizzini's mind* the types of Roberts have nothing to do with immunity.

Thus, altogether we should consider *2 different sets of types* of Roberts. In the following specification of types in $\ddot{T}_1$, the subscript represents the content of the glasses in their original order: "1" represents poison, "0" represents no-poison. The left digit in the subscript refers to the glass initially in front of Roberts, the right digit refers to the glass in front of Vizzini. The superscript has the letter "a" if Roberts is aware of the fact that Vizzini could have switched the glasses, and the letter "u" if Roberts is unaware of this fact; it has the letter "i" if Roberts is *i*mmune to the poison, and the letter "v" if Roberts is *v*ulnerable to the poison. However, in $\mathring{T}_1$ the latter coordinate is missing (describing types of Roberts *in the mind of Vizzini*)

$$\begin{aligned} \ddot{T}_1 &= \{t_{10}^{av}, t_{01}^{av}, t_{11}^{av}, t_{10}^{uv}, t_{01}^{uv}, t_{11}^{uv}, t_{10}^{ai}, t_{01}^{ai}, t_{11}^{ai}, t_{10}^{ui}, t_{01}^{ui}, t_{11}^{ui}\} \\ \mathring{T}_1 &= \{t_{10}^{a}, t_{01}^{a}, t_{11}^{a}, t_{10}^{u}, t_{01}^{u}, t_{11}^{u}\} \end{aligned}$$

The beliefs of the different types are listed in the following tables. Each row in a table corresponds to a type, and each column corresponds to a function in K. The entries specify the probabilities that the corresponding types ascribe to the corresponding functions in K.

$\ddot{T}_1$	k_{+-}	k_{-+}	k_{--}	k_{++}
t_{10}^{av}	$\frac{1}{2}$	$\frac{1}{2}$	0	0
t_{01}^{av}	$\frac{1}{2}$	$\frac{1}{2}$	0	0
t_{11}^{av}	0	0	1	0
t_{10}^{uv}	0	1	0	0
t_{01}^{uv}	1	0	0	0
t_{11}^{uv}	0	0	1	0
t_{10}^{ai}	$\frac{1}{2}$	0	0	$\frac{1}{2}$
t_{01}^{ai}	$\frac{1}{2}$	0	0	$\frac{1}{2}$
t_{11}^{ai}	1	0	0	0
t_{10}^{ui}	0	0	0	1
t_{01}^{ui}	1	0	0	0
t_{11}^{ui}	1	0	0	0

These beliefs reflect the idea that if Roberts is aware of the fact that Vizzini could have switched the glasses, Roberts assigns equal probabilities to the two possible cases – that the glasses have been switched or not switched.

$\mathring{T}_1$	k_{+-}	k_{-+}	k_{--}	k_{++}
t_{10}^{a}	$\frac{1}{2}$	$\frac{1}{2}$	0	0
t_{01}^{a}	$\frac{1}{2}$	$\frac{1}{2}$	0	0
t_{11}^{a}	0	0	1	0
t_{10}^{u}	1	0	0	0
t_{01}^{u}	0	1	0	0
t_{11}^{u}	0	0	1	0

Correspondingly, there are two spaces –

$$\begin{aligned} \ddot{\Omega} &= K \times \ddot{T}_1 \\ \mathring{\Omega} &= K \times \mathring{T}_1 \end{aligned}$$

which form a lattice – $\mathring{\Omega}$ is a quotient space of $\ddot{\Omega}$. This lattice of spaces constitutes an *unawareness structure* – a generalized kind of a type space introduced by Heifetz, Meier and Schipper (2005, forthcoming) for modeling *mutual* awareness and unawareness.

In this particular example the lattice of spaces is very simple, because only Roberts acts and has types in this unawareness structure. If there had been several players acting simultaneously, the lattice of spaces would have been more involved. Anyway, this example is meant to demonstrate that when *unawareness about an aspect of reality* is involved, game-states will constitute an *unawareness structure* rather than a classical type space.

Notice that if it were only for Roberts, we wouldn't need $\mathring{\Omega}$ in the lattice. However, we now turn to describe Vizzini's state of mind when Roberts has just turned around, and Vizzini has to decide whether or not to switch the glasses. In Vizzini's mind, each of these actions would lead to a state of affairs described by a game-state in $\mathring{\Omega}$. However, game-states in $\mathring{\Omega}$ do not represent actual game situations – they do not capture all the aspects relevant to the strategic interaction. The truth of the matter is that Vizzini's action leads to a game-state in $\ddot{\Omega}$. That's why we need the entire unawareness structure consisting of the two spaces $\ddot{\Omega}$ and $\mathring{\Omega}$.

In Vizzini's mind there are 6 functions that could potentially determine the game-state as a result of Vizzini's action – *N*ot to switch or to *S*witch

$$A_2' = \{N, S\}$$

These functions

$$\check{K}' = \{k_{10}^a, k_{01}^a, k_{11}^a, k_{10}^u, k_{01}^u, k_{11}^u\}$$

are from A_2' to $\mathring{\Omega}$: as a function of Vizzini's choice – N or S – they determine the payoff function in K and Roberts' type in $\mathring{T}_1$, as follows:

$$\begin{aligned}
k_{10}^a(N) &= (k_{+-}, t_{10}^a), \quad k_{10}^a(S) = (k_{-+}, t_{10}^a) \\
k_{01}^a(N) &= (k_{-+}, t_{01}^a), \quad k_{01}^a(S) = (k_{+-}, t_{01}^a) \\
k_{11}^a(N) &= (k_{--}, t_{11}^a), \quad k_{11}^a(S) = (k_{--}, t_{11}^a) \\
k_{10}^u(N) &= (k_{+-}, t_{10}^u), \quad k_{10}^u(S) = (k_{-+}, t_{10}^u) \\
k_{01}^u(N) &= (k_{-+}, t_{01}^u), \quad k_{01}^u(S) = (k_{+-}, t_{01}^u) \\
k_{11}^u(N) &= (k_{--}, t_{11}^u), \quad k_{11}^u(S) = (k_{--}, t_{11}^u)
\end{aligned}$$

We now have to describe Vizzini's types. In fact, Vizzini has a single type –

$$T_2' = \{t_2'\}$$

with the belief

$$b_2(t_2')(k_{10}^a) = b_2(t_2')(k_{01}^a) = b_2(t_2')(k_{10}^u) = b_2(t_2')(k_{01}^u) = \frac{1}{4}$$

Thus, Vizzini is uncertain whether his glass or Roberts' glass is poisoned, and ascribes equal weights to both possibilities; and, independently of this, Vizzini assigns equal probabilities to the possibilities that Roberts is aware or unaware of Vizzini's capability to switch the glasses.

The truth of the matter is that reality is described in richer terms than in Vizzini's mind, by the functions in

$$\hat{K}' = \left\{k_{10}^{av}, k_{01}^{av}, k_{11}^{av}, k_{10}^{ai}, k_{01}^{ai}, k_{11}^{ai}, k_{10}^{uv}, k_{01}^{uv}, k_{11}^{uv}, k_{10}^{ui}, k_{01}^{ui}, k_{11}^{ui}\right\}$$

from A_2' to game-states in $\ddot{\Omega}$. They are defined as follows:

$$\begin{aligned}
k_{10}^{av}(N) &= (k_{-+}, t_{10}^{av}), \quad k_{10}^{av}(S) = (k_{+-}, t_{10}^{av}) \\
k_{01}^{av}(N) &= (k_{+-}, t_{01}^{av}), \quad k_{01}^{av}(S) = (k_{-+}, t_{01}^{av}) \\
k_{11}^{av}(N) &= (k_{--}, t_{11}^{av}), \quad k_{11}^{av}(S) = (k_{+-}, t_{10}^{av}) \\
k_{10}^{ai}(N) &= \left(k_{++}, t_{10}^{ai}\right), \quad k_{10}^{ai}(S) = \left(k_{+-}, t_{10}^{ai}\right) \\
k_{01}^{ai}(N) &= \left(k_{+-}, t_{01}^{ai}\right), \quad k_{01}^{ai}(S) = \left(k_{++}, t_{01}^{ai}\right) \\
k_{11}^{ai}(N) &= \left(k_{+-}, t_{11}^{ai}\right), \quad k_{11}^{ai}(S) = \left(k_{+-}, t_{11}^{ai}\right) \qquad (*)
\end{aligned}$$

$$\begin{aligned}
k_{10}^{uv}(N) &= (k_{-+}, t_{10}^{uv}), \quad k_{10}^{uv}(S) = (k_{+-}, t_{10}^{uv}) \\
k_{01}^{uv}(N) &= (k_{+-}, t_{01}^{uv}), \quad k_{01}^{uv}(S) = (k_{-+}, t_{01}^{uv}) \\
k_{11}^{uv}(N) &= (k_{--}, t_{11}^{uv}), \quad k_{11}^{uv}(S) = (k_{+-}, t_{10}^{uv}) \\
k_{10}^{ui}(N) &= \left(k_{++}, t_{10}^{ui}\right), \quad k_{10}^{ui}(S) = \left(k_{+-}, t_{10}^{ui}\right) \\
k_{01}^{ui}(N) &= \left(k_{+-}, t_{01}^{ui}\right), \quad k_{01}^{ui}(S) = \left(k_{++}, t_{01}^{ui}\right) \\
k_{11}^{ui}(N) &= \left(k_{+-}, t_{11}^{ui}\right), \quad k_{11}^{ui}(S) = \left(k_{+-}, t_{11}^{ui}\right) \qquad (**)
\end{aligned}$$

Altogether, we now have an unawareness structure consisting of a lattice of two spaces –

$$\begin{aligned}
\hat{\Omega}' &= \hat{K}' \times T_2' \\
\check{\Omega}' &= \check{K}' \times T_2'
\end{aligned}$$

where the latter is a quotient space of the former. The situation is described by one of the game-states in $\hat{\Omega}'$ (in fact, in the film it is not clear whether it is (k_{11}^{ai}, t_2') or (k_{11}^{ui}, t_2')

– i.e. whether or not Roberts is aware that Vizzini could have switched the glasses). However, in Vizzini's mind the situation is described by a game-state in $\breve{\Omega}'$ (and with his type t_2' Vizzini ascribes equal probabilities to 4 of the game-states in $\breve{\Omega}'$).

We can now compute, by backward induction, the optimal action for each of the types. The optimal actions for the types of Roberts in $\ddot{T}_1, \mathring{T}_1$ are easy to compute. In particular, let us specify first the optimal action $\sigma_1(\cdot)$ of the following types in $\mathring{T}_1$:

$$\begin{aligned}
\sigma_1(t_{10}^a) &= rr \\
\sigma_1(t_{01}^a) &= rr \\
\sigma_1(t_{10}^u) &= rr \\
\sigma_1(t_{01}^u) &= dd
\end{aligned}$$

Why are these the optimal actions? The types t_{10}^a and t_{01}^a of Roberts are unsure whether the glasses have been switched, so the optimal action for them is to refrain from drinking: Drinking implies death of Roberts (payoff -2) with probability $\frac{1}{2}$ and death of Vizzini (payoff 1) with probability $\frac{1}{2}$. (Recall that these types are in the mind of Vizzini, so immunity is unimaginable!) The expected payoff from drinking is thus negative, while refraining from drinking yields a sure payoff of 0. In contrast, the unaware types t_{10}^u and t_{01}^u act on the premise that the glasses are in their original position, so of these two types only t_{01}^u (who is certain that only Vizzini's glass is poisoned) chooses to drink.

This means that if we fold the game-states in $\mathring{\Omega}$ with this optimal strategy of Roberts, then

$$\begin{array}{lll}
(k_{+-}, t_{10}^a) & \text{is replaced by the payoff vector } (0,0) & \text{(a)} \\
(k_{-+}, t_{10}^a) & \text{is replaced by the payoff vector } (0,0) & \text{(b)} \\
(k_{-+}, t_{01}^a) & \text{is replaced by the payoff vector } (0,0) & \text{(c)} \\
(k_{+-}, t_{01}^a) & \text{is replaced by the payoff vector } (0,0) & \text{(d)} \\
(k_{-+}, t_{10}^u) & \text{is replaced by the payoff vector } (0,0) & \text{(e)} \\
(k_{-+}, t_{10}^u) & \text{is replaced by the payoff vector } (0,0) & \text{(f)} \\
(k_{+-}, t_{01}^u) & \text{is replaced by the payoff vector } (1,-2) & \text{(g)} \\
(k_{-+}, t_{01}^u) & \text{is replaced by the payoff vector } (-2,1) & \text{(h)}
\end{array}$$

With his type t_2', Vizzini believes that choosing Not to switch the glasses will lead, with equal probabilities, to one of the 4 game-states mentioned in (a),(c),(e),(g), and hence to a negative expected payoff:

$$\frac{1}{4} \times 0 + \frac{1}{4} \times 0 + \frac{1}{4} \times 0 + \frac{1}{4} \times (-2) = -\frac{1}{2}$$

while, according to this belief, choosing to Switch the glasses will lead, with equal probabilities, to one of the game-states mentioned in (b),(d),(f),(h), and hence to a positive expected payoff to himself:

$$\frac{1}{4} \times 0 + \frac{1}{4} \times 0 + \frac{1}{4} \times 0 + \frac{1}{4} \times 1 = \frac{1}{4}$$

So the optimal action of (the unique type t_2' of) Vizzini is S.

However, the truth of the matter is that both glasses are poisoned and Roberts is immune to the poison – the true game-state is either (k_{11}^{ai}, t_2') or (k_{11}^{ui}, t_2'). From $(*)$ and $(**)$ it follows that in either case, when Vizzini chooses S (and in fact even if he were to choose N) the resulting payoff function is k_{+-}, and the type of Roberts is either t_{11}^{ai} or t_{11}^{ui}. From lines 9 and 12, respectively, in the table defining the beliefs of the types in $\ddot{T}_1$, we see that with either of these types Roberts is certain that the payoff function is k_{+-}. From the definition of k_{+-} we see that the optimal action for Roberts is dd.

7 References

Aumann, R., S. Hart and M. Perry (1997), "The Absent-Minded Driver," *Games and Economic Behavior* **20**, 102-116.

Heifetz, A., M. Meier and B. Schipper (forthcoming), "Interactive Unawareness,", *Journal of Economic Theory*

Heifetz, A., M. Meier and B. Schipper (2005), "Unawareness, Beliefs and Games", mimeo

Heifetz, A. and A. Pauzner (2005), "Backward Induction with Players who Doubt Others' Faultlessness," *Mathematical Social Sciences* **50**, 252-267

Gilboa, I. (1997), "A Comment on the Absent Minded Driver Paradox," *Games and Economic Behavior* **20,** 25-30

Lipman, B. (1997), "More Absentmindedness," *Games and Economic Behavior* **20,** 97-101.

Piccione, M., and A. Rubinstein (1997), "On the Interpretation of Decision Problems with Imperfect Recall," *Games and Economic Behavior* **20**, 3-24.