

Nash equilibrium, rational expectations, and heterogeneous beliefs: Action-consistent Nash equilibrium

Jernej Čopič *

This draft: July 26, 2012; first draft: September 1, 2011.

Abstract

In an action-consistent Nash equilibrium of a simultaneous-moves game with uncertainty a la Harsanyi (1967) players choose optimally, make correct assessments regarding others' actions, and infer information from others' optimal play. In equilibrium, players may have differing information in the sense of prior beliefs. An example of action-consistent Nash equilibrium strategy profile is constructed, which is not a Nash equilibrium under any common prior. Given the equilibrium strategy profile, each type of each player faces pooling of the other players' types. In a modified example one type of one player doesn't face such a pooling problem: the strategy profile can still satisfy equilibrium

*jcopic@econ.ucla.edu. The idea for this work stems from the paper titled Awareness Equilibrium, which was written jointly with Andrea Galeotti, and is available as a working paper. I am especially grateful to Pierpaolo Battigalli, and also to Andy Atkeson, Andrea Galeotti, Matthew O. Jackson, Debrah Meloso, Joe Ostroy, Mike Rothschild, Hugo Sonnenschein, Bill Zame, seminar audiences at Bocconi, UCLA econometrics proseminar, USC, University of Amsterdam, University of Ljubljana, Dept. of Mathematics, Analytical Research Center at the BANK of slovenia, and participants of the conference in honor of 20 years of IDEA. Jeeheh Choi's valuable research assistance is also gratefully acknowledged. I also wish to thank Pablo Branas and Filippos Exadaktylos at the University of Granada for their hospitality.

conditions up to an arbitrary finite order, but no longer in an infinite regress of action-consistent Nash equilibrium.

1 Introduction

Think of two sophisticated people, call them 1 and 2, who encounter each other in a strategic situation with uncertainty. Each of them privately observes some relevant information regarding the uncertain parameters and must decide what to do based on that information. The information that 2 might have could be relevant to 1's payoff in itself. Player 1 will also try to guess what information 2 has from 2's choice of actions – Player 1 will thus trim her guesses regarding 2's behavior by assessing whether a rational opponent would be prone to follow a particular course of action. It seems that people often engage in such strategic analysis and failure to do so may have serious consequences – a rational and sophisticated player should not fall prey to such mistakes.^{1,2} In equilibrium, players choose optimally, and the assessments they make are consistent with what each of them sees and with opponent's rationality – they infer information from opponent's optimal play. In equilibrium, players may have differing information, but no player has any reason to change either her strategy or her assessments.

The idea that economic agents' optimal choices and their beliefs should be determined simultaneously is at the heart of almost any definition of economic equilibrium. Agents hold beliefs regarding economic environment and make choices; in equilibrium

¹Duke Wellington presumably engaged in such strategic analysis: “All the business of war, and indeed all the business of life, is to endeavour to find out what you don't know by what you do; that's what I called ‘guessing what was at the other side of the hill.’” http://en.wikiquote.org/wiki/Arthur_Wellesley,_1st_Duke_of_Wellington.

²A famous Danish chess player from early 20th century, Aaron Nimzowitsch apparently burst out on an occasion of having lost an important game: “Gegen diesen Idioten muss ich verlieren!” (“That I should lose to this idiot!”) – http://en.wikipedia.org/wiki/Aron_Nimzowitsch#Personality. While the object of study here are simultaneous-moves games, such underestimation of one's opponent would be a particular violation of equilibrium requirements described here.

their beliefs are correct and their choices optimal. For example, in a Nash equilibrium under complete information (c.f. Nash (1951)), agents optimally react to what they believe the others might be doing and different such correct beliefs will give rise to different equilibria. This idea is particularly evident in games with discontinuous payoffs, where equilibrium must be given such an interpretation in order for equilibrium to exist at all: beliefs regarding randomization must be (correctly) specified in equilibrium, see e.g., Simon and Zame (1990). In the setting of a general equilibrium with uncertainty, in a Rational-expectations equilibrium (c.f. Radner (1982)) agents hold an endogenously determined belief regarding equilibrium prices: in equilibrium their choices are optimal, their beliefs are correct, and information is inferred from prices. The setting of this paper is a simultaneous-moves game with uncertainty and this equilibrium idea is formulated as action-consistent Nash equilibrium.

Harsanyi (1967) formalized a simultaneous-moves game with uncertainty by showing that all uncertainty regarding fundamentals can be reduced to uncertainty regarding the players' payoffs. Harsanyi introduced a move by *Nature*, whereby a type is drawn for each player from a joint distribution, and the vector of types then defines the players' payoffs in the game. Harsanyi defined a *C* – *game* as one where the belief of each type of each player can be derived from a common prior, which is the (objective) distribution of Nature's moves. In a *C* – *game*, Harsanyi (1968 a,b) defined a Nash equilibrium under a common prior: each player chooses an optimal action for each of her types, in order to maximize her expected payoff given her belief, and given such type-contingent strategy of all other players.

Implicit in a Nash equilibrium of a *C* – *game* is a notion of a steady state where each player has learned everything that could be learnt: each player has correct probabilistic information regarding the likelihoods of all type draws. Fudenberg and Levine (1993b) provide a formal learning foundation for a more general class of steady states, called Self-confirming equilibria. They envision a recurrent play of the game by a large population of players where repeated-game effects are absent as players are

unlikely to meet the same opponent twice.³ The notion of Self-confirming equilibrium is more general than that of a Nash equilibrium in a C – game in that each player only learns some privately-observable facts, described by a feedback function (also called a signal function). Hence, in a Self-confirming equilibrium players’ beliefs need not be consistent with the common prior, but rather just some aspects of the distribution of Nature’s draws and players’ behavior as described by the feedback function.⁴

Another difference between Harsanyi’s equilibrium and Self-confirming equilibrium is that in a Self-confirming equilibrium a player need not be able to justify the others’ behavior as a result of some optimizing procedure. In a Nash equilibrium of a C – game such *rationalization* comes for free: if a player is optimizing against the distribution over the opponents’ actions, she can in particular justify the other players’ behavior by their actual strategies, which are evidently optimal, given that all players have the same beliefs. Moreover, each player can in principle assign such rationalization of others’ behavior to all other players, and can *a fortiori* proceed with such rationalization in an infinite regress. Hence, in Harsanyi’s definition of equilibrium, players need not actually know the other players’ strategies – it is just *possible* that they might assign optimal strategies to their opponents. A player’s belief can then be thought of as being not only over types but also over the other players’ strategies – to avoid confusion such more general belief is here called *an assessment*. In Harsanyi’s definition of Nash equilibrium in a C – game it is thus implicit that each player’s strategy survives a rationalizability criterion from the per-

³Harsanyi envisioned the common prior as resulting from some unspecified common learning experience. In the spirit of Fudenberg and Levine (1993b) one can imagine a specific such learning process where players’ types are drawn each time anew, and the information regarding all past type draws is available in some repository accessible to all players. When each player privately accesses the repository without the other players knowing that, all players will eventually hold a common prior but not more can be said about the players’ higher-order knowledge. For a characterization of Nash equilibrium in terms of players’ knowledge see Aumann and Brandenburger (1995).

⁴For Self-confirming equilibrium see also Fudenberg and Levine (1993a) and Dekel et al. (2004). Self-confirming equilibrium is very closely related to Conjectural equilibrium, see Battigalli (1987), Rubinstein and Wolinsky (1994), Battigalli and Guatoli (1997), and Esponda (2011).

spective of all other players, while players' beliefs coincide with the distribution of Nature's moves.⁵

Here, Harsanyi's definition of Nash equilibrium in a C – *game* is extended to a more general class of simultaneous-moves games with uncertainty with specific kind of non-common priors. The equilibrium notion here is related to Copic and Galeotti (2007), which effectively uses the epistemic structure from Bernheim (1984), along with a notion of consistency of beliefs. The idea of rationalizability where a steady state results from deductive learning is thus combined with the idea of self-confirming equilibrium where a steady state results from inductive learning of some privately observable parameters as specified by the feedback function. Here, this feedback function is given by a player's own payoffs and the distribution over others' actions, conditional on the player's type and actions that this type chooses to play; each player also knows the marginal distribution over her own types. The name *action-consistent* Nash equilibrium refers to the particular feedback function. The equilibrium notion then strenghtens the optimizing requirement of Self-confirming equilibrium to the aforementioned rationalizability requirement. While no formal learning foundation is provided here, *once in* an action-consistent Nash equilibrium, each player would have no reason to change either her belief regarding uncertain parameters, her assessment regarding the others' play, or her behavior.⁶

Our formal definition of equilibrium is that players' strategies are supported by conjectures satisfying *a common belief in optimization and action-consistency*. A common belief in optimization and action-consistency means, for example, that player 1 only needs to be able to attribute optimizing behavior to player 2 with

⁵The notion of rationalizability is posterior to Harsanyi's work. For complete-information versions of rationalizability see Bernheim (1984) and Pearce (1984). For a comprehensive treatment of incomplete-information versions of rationalizability see Battigalli et al. (2009).

⁶A heuristic learning interpretation of the equilibrium conditions is analogous to that in Fudenberg and Levine (1993b) and Dekel et al. (2004). The difference is that learning would here presumably have to also be deductive: in every recurrent play of the game a player would also have to learn from rationalizing the opponents' play.

respect to some belief that player 2 might have given player 1's information, not with respect to the belief that player 1 has given her own information. Intuitively, in equilibrium players cannot reject the possibility that all players are consistent, rational, and sophisticated.⁷

Our paper is focused on an example to illustrate how this relaxation of Harsanyi's common-prior requirements explains behavior which cannot be explained by assuming a common prior. To do that formally, another idea is borrowed from Harsanyi – that of a properly-informed observer who does not observe more than the players do. Since in an action-consistent Nash equilibrium players can be thought of as observing each other's actions and only own types and payoffs, such properly-informed observer can observe the distribution over players' actions and not more than that – since the observer is properly informed, the actual distribution over Nature's moves and the actual players' strategies are not observable to him. The example is given by a game in which the observer can explain a particular distribution over players' actions in a Nash equilibrium under diverse priors, but this cannot be explained as a Nash equilibrium under any common prior.

The example here is the simplest possible: there are two players, each has two types, and each type of each player has two actions. Essential to the example is a

⁷Common belief in optimization and consistency is closely related to Brandenburger and Dekel (1987) and Battigalli and Siniscalchi (2003), who characterize Nash equilibrium as rationality and a common belief in rationality. In a recent paper, Esponda (2011) defines equilibrium as a common belief in rationalizability and a common belief in a consistency criterion, called *confirmation*. Confirmation is different from action-consistency. On a fundamental level, the difference is that action-consistency embraces Harsanyi's idea that all fundamental uncertainty can be represented as uncertainty regarding the players' payoffs. Esponda's *confirmation* is derived from Vickrey (1961), where uncertainty is represented by a draw of the state of Nature (or a state of "fundamentals" of the economy), and each player receives a signal regarding that state. Confirmation requires that players' beliefs be consistent with the *actual* draw of the state of Nature, according to a given feedback function and the noise structure of each player's signal. This is quite different from action-consistency derived from Harsanyi's definition – given the same feedback function, action-consistency is stronger. It should also be noted that while most of Esponda's paper refers to a complete-information case, he defines his equilibrium notion in a more general belief-type space, and introduces all the necessary machinery to do so. Esponda's main objective is a characterization of his equilibrium notion as iterative deletion of actions that are not a best response in a particular sense.

particular kind of pooling: each type of a player cannot perfectly identify the conditional distribution over the other's types, conditional on the action that she is taking, but she could identify this conditional distribution had she taken the other action. This is true for each type of each player. Informally, the example is canonical in that it is given by the minimal necessary structure for such a situation to arise – that is, a situation, which cannot be explained as a Nash equilibrium under any common prior but can be explained as an action-consistent Nash equilibrium. By modifying the example slightly to one where only one type of one player does not face such pooling problem, the corresponding strategies are no longer equilibrium strategies under non-common priors. However, in this modified example the objective distribution of Nature's moves can be chosen so that the corresponding strategies satisfy equilibrium conditions up to an arbitrary finite order. This provides an illustration of the difference between requiring a common belief in optimization and consistency in an infinite regress, and requiring only a finite-order of such common belief.

The level of generality and notation introduced here does not go beyond what is necessary to carefully formulate our example. In our definition, for example, players do not attribute any correlation to other players' strategies.⁸ Players also make point-assessments regarding other players' strategies, and each type of each player holds the same belief.^{9,10} Effectively then, our example holds under assumptions that

⁸When players accumulate data in a recurrent play of a game correlation between players strategies could in general arise. Moreover players could have private assessments regarding such correlation, see for example Aumann and Brandenburger (1995). In an extensive-form setting, Fudenberg and Kamada (2011) provide several examples illustrating this issue of correlation.

⁹If one considers rationalizability alone, then there is a difference between allowing probabilistic assessments and point-wise assessments even in the case of complete information, see Bernheim (1984). However, additional requirement of consistency of beliefs provides a further restriction on these assessments, and in particular, point-wise strategy assessments are then without loss of generality.

¹⁰In a general setting, players' type space is given by a belief-type space, formally constructed by Mertens and Zamir [1985]. However, in our particular setting, restricting to beliefs that are only over payoff types is without loss of generality: if each type has a different belief over general belief types, and such belief has to be consistent with the same data, then in particular, all types can hold the same belief. Thus, the relevant belief-type subspace is effectively isomorphic to an infinite regress of beliefs over payoff types. We leave this justification at a heuristic level as formalizing it

are stronger than necessary while still consistent with extending Harsanyi's definition of Nash equilibrium in a C – game.

In Section 2, the basic model is given. In Section 3, the notion of a properly-informed observer for the present setting is defined. In Section 4 there is the main example. In Section 5, there is the example illustrating the difference between action-consistent Nash equilibrium and requiring only a finite-order belief in optimization and consistency.

2 Model

For a set X , denote by $\Delta(X)$ the set of lotteries, or probability distributions, over X . When X is a subset of an Euclidean space, and $\mu \in \Delta(X)$, then E_μ denotes the expectation operator with respect to lottery μ . For $\mu \in \Delta(X)$, denote by $\mu(x)$ the probability of $x \in X$. When $X = X_1 \times X_2$, denote by $\mu_{X_2|X_1'}$ the conditional distribution over X_2 given $X_1' \subset X_1$, and by μ_{X_1} marginal distribution over X_1 ; we will sometimes denote by $\mu_{x_2|x_1}$ the probability of $x_2 \in X_2$ conditional on $x_1 \in X_1$. Finally, 1_X denotes the indicator function, i.e., taking value 1 on set X and 0 everywhere else.

SIMULTANEOUS-MOVES GAME WITH UNCERTAINTY, PRIORS, RISK, STRATEGIES.

A simultaneous-moves game with uncertainty is defined as $\Gamma = \{N, A, T, u\}$, where $N = \{1, \dots, n\}$ is a finite set of players, $A = \times_{i \in N} A_i$ is a product of finite sets of players' actions, $T = \times_{i \in N} T_i$ is a product of finite sets of players' types, and $u : A \times T \rightarrow R^{|N|}$ is a vector of players' payoff functions. Throughout it is assumed that Γ is common knowledge: every player knows Γ , knows that every other player knows Γ , and so on. Types are drawn by Nature according to a (objective) probability

would introduce a large amount of additional notation; one can also consider this to be another particular restriction of our example. Esponda (2011) gives a formal proof of a similar statement in a different setting, as described above.

distribution $P^0 \in \Delta(T)$, where $P^0(t)$ is the probability of a type draw $t \in T$. We assume that the marginal probability of each type of each player is positive, $P_i^0(t_i) > 0, \forall i \in N, \forall t_i \in T_i$. A game with risk is given by (Γ, P^0) . Here it is not assumed that P^0 be common knowledge; it is assumed that P^0 is *objective*, in the sense that if an observer had an ideal set of observations of Nature's moves, he could identify P^0 . Player i 's strategy (mixed or pure) is a mapping $s_i : T_i \rightarrow \Delta(A_i)$, so that $s_i[t_i] \in \Delta(A_i)$ is the randomization over player i 's actions, when i 's type is t_i . A strategy profile is given by $s = (s_1, s_2, \dots, s_n)$. When players' types are $t \in T$ and their strategy profile s , the vector of players' payoffs is denoted by $u(s, t)$, in a slight abuse of notation.

A standard situation is one in which prior P^0 is common, that is, all players have the same belief about the probability distribution of Nature's moves. Following Harsanyi, a Nash equilibrium of (Γ, P^0) with a common prior P^0 is then usually defined as a strategy profile s^* , such that for each of player i 's types, his randomization over actions achieves the highest expected payoff given the conditional probability distribution over the other players' types and their strategies.

NASH EQUILIBRIUM UNDER A COMMON PRIOR. A strategy profile s^* is a Nash equilibrium (NE) strategy profile of (Γ, P^0) under a common prior P^0 if,

$$a_i^* = \arg \max_{a_i \in A_i} E_{P_{-i|t_i}^0} E_{s_{-i}[t_{-i}]} u_i(t_i, t_{-i}, a_i, a_{-i}), \quad \forall a_i^* : s_i|t_i(a_i^*) > 0, \quad \forall t_i \in T_i.$$

It is useful to single out *optimality for i* as the property that i 's strategy s_i be optimal with respect to a probability distribution P , given the other players' strategies s_{-i} . Optimality for i is hence a property of P and s .

OPTIMALITY. Let s be a strategy profile, let P be a prior belief and let $i \in N$.

Then, P and s_i satisfy *optimality for i* if,

$$a_i^* = \arg \max_{a_i \in A_i} E_{P_{-i|t_i}} E_{s_{-i}[t_{-i}]} u_i(t_i, t_{-i}, a_i, a_{-i}), \quad \forall a_i^* : s_{i|t_i}(a_i^*) > 0, \quad \forall t_i \in T_i. \quad (1)$$

ASSESSMENTS.¹¹ A central point pursued here is that players do not necessarily have true beliefs regarding the distribution over the states of the world and over the other players' play – players form *assessments*. An assessment by player $i \in N$ is a pair P, s , where $P \in \Delta(T)$ is a belief assessment, and $s \in \times_{j \in N} \Delta(A_j)$ is a strategy assessment.¹²

To close the model, players' assessments are anchored by the requirement that whatever data a player can observe should be consistent (in a probabilistic sense) with the actual distributions generating these data, i.e., with P^0 and s . Here it is assumed that each player can observe, the marginal distribution over her own types, and conditional on own type and realized action, the joint distribution over own payoffs and the other players' actions.¹³ Observing own payoffs is a natural requirement as it is hard to imagine what it means to play a *one-shot* game without knowing what payoff was obtained at the end. Hence, the key part of the observability assumption is that apart from own type, and knowing own strategy, each player observes the distribution over the opponents' actions. Thus, this consistency requirement imposed

¹¹The word “assessment” instead of “belief” is used here to distinguish from the more common use where beliefs are over types alone. An assessment here thus has two components: one describing a belief assessment over the joint distribution of types, and two, describing the point-assessment of the strategy profile being played.

¹²Strategy assessments here are point assessments. This is much simpler, and is here without loss of generality in the following sense: for any probabilistic equilibrium strategy assessment (appropriately defined) there is a deterministic assessment supporting the same outcome.

¹³In Harsanyi's description of *C – games* the model is also closed, as players have a common prior, and their common beliefs can be thought of as being entirely anchored by the truth. Harsanyi thus restricts the definition of equilibrium to *C – games*, and refrains from defining an equilibrium in *I – games*. For an early discussion of common-prior assumption, where players are allowed to have any beliefs whatsoever see Morris (1995).

on assessments is called *action-consistency for i* .¹⁴

A bit of additional notation is introduced to define action-consistency. Let $V_i \subset \mathbb{R}$ be the set of payoffs which are attainable by player i in game Γ , i.e., $V_i \equiv \text{Image}(u_i)$. Since Γ is common knowledge, the product $V = \times_{i \in N} V_i$ is common knowledge. Let $\tilde{G}[P, s]$ be the distribution over $T \times A \times V$, resulting from a prior P and a strategy profile s . That is, $\tilde{G}[P, s]$ is given by a matrix with elements $\tilde{g}[P, s]_{t,a,v}$, $t \in T, a \in A, v \in V$, where,

$$\tilde{g}[P, s]_{t,a,v} = P(t) \times s_1[t_1](a_1) \times \dots \times s_1[t_1](a_1) \times 1_{\{v=u(t,a)\}}.$$

ACTION-CONSISTENCY FOR i . Let P be a probability distribution over types, s a strategy profile, and let $i \in N$. An assessment P^i, s^i is *action-consistent with P and s for i* , if, $s^i \equiv s_i, P_{T_i}^i \equiv P_{T_i}$, and,

$$\tilde{G}[P^i, s^i]_{T_i \times A \times V_i | t_i, a_i} \equiv \tilde{G}[P, s]_{T_i \times A \times V_i | t_i, a_i}, \forall t_i \in T_i, \forall a_i \in A_i, \text{ s.t., } s_i[t_i](a_i) > 0. \quad (2)$$

To illustrate action-consistency, consider the following example.¹⁵

EXAMPLE 1. There are two players, $N = \{1, 2\}$, their actions are given by $A_1 = \{up, down\}, A_2 = \{L, R\}$, their types are $T_1 = \{t_1\}, T_2 = \{t_2, t'_2\}$, and the distribution over these types is

$P^0(t_1, t_2) = \frac{3}{4}, P^0(t_1, t'_2) = \frac{1}{4}$. The payoff structure is given by the following two

¹⁴Players' assessments need not be consistent with the *whole outcome* of the game, that is, the joint distribution of types and actions, and hence also every player's payoffs. If that were the case, then such stronger consistency would imply that players must have a common prior equal to the distribution of Nature's moves. Under action-consistency players' beliefs can be consistent and nonetheless different. If consistency is interpreted as what players observe in a recurrent play of the game, then observing the whole outcome might sometimes not be possible, when action-consistency may be feasible. See Fudenberg and Levine (1993b) and Dekel et al. (2004) for learning in recurrent games.

¹⁵This example is motivated by an example in Jackson and Kalai (1997).

payoff matrices.

(t_1, t_2)	L	R
up	$\alpha, 1$	$0, -1$
down	$-10, 0$	$0, -1$

(t_1, t'_2)	L	R
up	$0, 1$	$0, -1$
down	$4, 0$	$0, -1$

Suppose that $\alpha = 0$, and let $s^0 = (up; R \mid t_2, R \mid t'_2)$. In that case player 1 can hold a variety of different action-consistent assessments regarding the probability of (t_1, t_2) and the strategy profile s^0 – in fact, in order for her assessment to be action consistent, she can hold any assessment. Similarly, if $s^0 = (up; L \mid t_2, L \mid t'_2)$ and $\alpha = 0$. However, when $s^0 = (up; L \mid t_2, L \mid)$ and $\alpha \neq 0$, only the assessment P^0, s^0 is action consistent with P^0, s^0 . Finally, note also that when $s^0 = (up; L \mid t_2, L \mid)$ and $\alpha = 0$, optimality imposes an additional constraint: in order for player 1 to optimize while playing *up*, she must hold a belief which puts at most $\frac{4}{14}$ on type (t_1, t_2) .

Another interpretation of optimality and action-consistency is quite familiar. Consider an extensive-form game where Nature moves first, each player observing only her own type, and players then move simultaneously. This is precisely the game of incomplete information as defined by Harsanyi (1967-1968). Action-consistency is then the requirement that, given player i 's information sets, all final nodes *that are actually reached and that player i can identify*, are reached with correct probabilities. In this extensive form, player i has complete information about her own type and action that she took, so that player i should be able to identify the partition of final nodes as given by $A_{-i} \times V_i$, and the events in this partition have to be reached with correct probabilities. Hence, action-consistency as defined here is simply consistency in Harsanyi's interpretation of a simultaneous-moves game of incomplete information when players do not have a completely common learning experience – they do not observe each others' types or payoffs. ¹⁶

¹⁶In the work on self-confirming equilibrium and conjectural equilibrium, such partitional infor-

The final ingredient of the present model is the highest possible level of rationality, or sophistication, of players playing the game. Not only are players consistent and optimize, but whatever they perceive regarding the distribution of Nature's draws and the others' play, they also *should be able to justify* that as a play by similarly sophisticated players. In order to conceptualize such considerations, higher-order assessments are defined, and each player holds a *hierarchy* of assessments. Such hierarchy of assessments is called a *conjecture*.¹⁷

For $k \in \mathbb{N}$, let N^k be the set of all sequences of length k with elements from N ; for $k = 0$, define $N^0 = \emptyset$. Let \mathcal{L}^i be the set of all finite sequences of integers from N , such that the first element is i , and no two adjacent elements are the same.¹⁸

That is, $\mathcal{L}^i = \{\ell \in N^k; k \geq 1, \ell_1 = i, \ell_{m+1} \neq \ell_m, 1 \leq m \leq k\}$. Also, let $\mathcal{L} = \{\ell \in N^k; k \geq 1, \ell_{m+1} \neq \ell_m, 1 \leq m \leq k\}$, so that $\mathcal{L} = \cup_{i \in N} \mathcal{L}^i$. For sequences $\ell, \ell' \in \mathcal{L}$, and an integer $i \in N$, define the following operations: $l(\ell)$ is the last element of ℓ ; (ℓ, i, ℓ') is given by elements from ℓ , followed by i , and followed by the elements from ℓ' ; (ℓ, \emptyset) is the sequence ℓ .

mation structure is formally given by a feedback function. Fudenberg and Kamada (2011) study more general *partition-confirmed equilibrium* in extensive-form games, where Nature's moves are known to all players, but players might be unable to observe each other's moves at various information sets.

¹⁷As mentioned in the introduction, only payoff types are considered here; all beliefs and assessments are specified over payoff types only. The main objective of this paper is to provide a specific example of an equilibrium outcome of a game, played by players who are as rational as they can be and have heterogeneous beliefs. Given that the outcome constructed here is supportable already by assessments that are over payoff types only, this outcome would constitute an equilibrium outcome under action consistency, if more general belief types were considered, as defined by Mertens and Zamir (1985). Working with payoff types alone is much simpler and it suffices for the purpose.

¹⁸This is a very similar construction and notation to that in Bernheim (1984) – the only difference is that the first element of the sequence here describes the player herself, whereas in Bernheim's (1984) notation the first element of the sequence describes players assessment about some other player. Here too, the assessments are limited to be only about other players, and not about the player herself. This is here without loss of generality. Action-consistency would impose strong constraints on a player's assessments about herself, and moreover, any such equilibrium would have a strategically equivalent equilibrium where a player's assessment about her own assessment coincides with her actual assessment.

CONJECTURES. A conjecture by player $i \in N$ is a hierarchy of belief and strategy assessments,

$$C_i = (P^\ell, s^\ell)_{\ell \in \mathcal{L}^i}.$$

P^ℓ and s^ℓ , $\ell \in \mathcal{L}^i$, are interpreted as some order of assessments of player i about player j , where $j = l(\ell)$.¹⁹ For a conjecture C_i , denote by C_i^ℓ the ℓ -assessment of player i . A profile of conjectures is given by $C = (C_1, C_2, \dots, C_n)$; when $\ell = (i)$, $C_i^{(i)}$ is the assessment of player i about Nature's moves and the actual strategy profile being played.

A common belief in optimality and a common belief in action-consistency are defined inductively.

COMMON BELIEF IN OPTIMALITY. Let C be a profile of conjectures, and for $i \in N$, let $(P^\ell, s^\ell) = C_i^\ell$, $\ell \in \mathcal{L}^i$. C satisfies a common belief in optimality if P^ℓ and s_j^ℓ satisfy optimality for $j = l(\ell)$, for every $\ell \in \mathcal{L}^i$, for every $i \in N$.

COMMON BELIEF IN ACTION-CONSISTENCY. Let P^0 be a probability distribution of Nature's moves, and s^0 a strategy profile. Let C be a profile of conjectures, and let $(P^\ell, s^\ell) = C_i^\ell$, for each $i \in N$ and each $\ell \in \mathcal{L}^i$. C satisfies a common belief in action-consistency if,

1. $P^{(i)}$ and $s_i^{(i)}$ are action-consistent for i with P^0 and s^0 , for every $i \in N$;
2. $P^{(\ell, j)}$ and $s^{(\ell, j)}$ are action-consistent for j with P^ℓ and s^ℓ , for every $j \in N$, for every $\ell \in \mathcal{L}^i$, for every $i \in N$.

DEFINITION 1. Let P^0 be the probability distribution over Nature's moves, let C^* be a profile of conjectures, and let s^* be a strategy profile. C^* is an action-consistent

¹⁹Order of an assessment is here the length of ℓ . Note that when $n > 2$, there are many different assessments of the same order.

equilibrium of (Γ, P^0) for P^0 and s^* , if C^* satisfies a common belief in optimality and action-consistency with P^0 and s^* .

Then, s^* is then supportable in an action-consistent Nash equilibrium of (Γ, P^0) .

As an illustration why common belief in consistency and optimality as a definition of equilibrium is reasonable, consider the following example. This example illustrates implicit assumptions regarding players' sophistication when they only act optimally and consistently. Hence, this type of "Nimzowitsch mistake" (see footnote 2), when players do not make additional inference based on each other's rationality, is in itself a particular failure of rationality. By failing to take into account the incentives of her opponent, a player may incur large losses – these losses may be arbitrarily large even while the incentives of the opponent may be arbitrarily strong.

EXAMPLE 2. Let player 1 have only one type, $T_1 = \{t_1\}$, while 2 has two types, $T_2 = \{t_2, t'_2\}$. Each player has two actions, $A_1 = \{u, d\}$, $A_2 = \{L, R\}$, and the payoff structure for each draw of types is specified as follows.

(t_1, t_2)	L	R
u	-100, 100	-100, 0
d	1000, 0	-100, 1

(t_1, t'_2)	L	R
u	-100, 0	-100, 100
d	-10000, 0	1001, 1000

The distribution of Nature's draws P^0 is given by,

P^0	t_2	t'_2
t_1	$\frac{1}{2}$	$\frac{1}{2}$

Consider player 1's belief given by $P^1 = P' = (\epsilon, 1 - \epsilon)$, where $\epsilon > 0$, and let $P^2 = P^0$. Consider a strategy profile $s = (u; L|t_2, R|t'_2)$. It is immediate that for both players the strategy profile s along with these beliefs satisfy optimality and consistency with P^0 and s in (Γ, P^0) . An interpretation of such profile s is that

player 1 fails to make any inference whatever from observation of player 2's play. When she plays u , player 1 cannot learn anything about P^0 from *only* observing the frequencies of her own payoffs. Under the assumption that player 1 observes player 2's play her assessment implies a presumption that player 2's behavior is irrational; under such a presumption player 2 could equally well hold any beliefs regarding the distribution over player 2's types.

For example, player 1 could then make the assessment of player 2's strategy $s^{(1,2)} = (u; R|t_2, \frac{1+2\epsilon}{2} \times L + \frac{1-2\epsilon}{2} \times R|t'_2)$, which satisfies action-consistency. Allowing for such assessment would be equivalent to assuming that player 1 has no knowledge about player 2's payoffs. The problem of preference revelation would be super-imposed over the strategic problem, which would implicitly contradict the assumption of common knowledge of Γ . Additionally, if player 1 is allowed to make such an assessment, then one should not have any qualms with player 1 further supposing that 2 could play an even stranger strategy if 1 were to play d , for example, $(R|t_2, L|t'_2)$, which should further convince player 1 to play u . That is, if player 1 can "in equilibrium" assume non-optimizing behavior for player 2, then she should be allowed to assume outright crazy behavior on counterfactual events, such as playing strictly dominated actions. This latter assumption would be testable for player 1 from variation in her own payoffs if she were to actually play d . Hence, allowing for a player to make strange inferences regarding the behavior of her opponent would invariably lead to various kinds of implicit inference assumptions, which might be hard to justify. Common knowledge of Γ , along with a common belief in optimality and action-consistency are natural internally consistent axioms for a simultaneous-moves game under uncertainty as defined by Harsanyi, under the assumption that players' actions are observable.

Action-consistency allows for heterogeneity in players' beliefs. It also guarantees that whenever for whatsoever reason players do have a common prior, an action-

consistent Nash equilibrium is a Nash equilibrium. Heuristically, if players have the same common prior equal to the true distribution of Nature's moves, and there is a common belief in action-consistency and optimality, then in particular, each player optimizes with respect to that common prior.

PROPOSITION 1. If a hierarchy of conjectures C^* is an action-consistent equilibrium of (Γ, P^0) for a probability distribution P^0 and strategy profile s^* , and such that under C^* all players have (first-order) belief P^0 about the distribution of Nature's moves, then s^* is a Nash equilibrium of (Γ, P^0) under the common prior P^0 .

PROOF. Let C^* be a profile of conjectures such that there is a common belief in optimality and consistency, and players have a common prior. That is, $C_i^{(i)*} = (P^{(i)}, s^{(i)})$, where $P^{(i)} = P^{(j)}$, $\forall i, j \in N$. By assumption, $P^{(i)*}$ and $s_i^{(i)}$ satisfy optimality for i , and $P^{(i)}$ and $s^{(i)}$ are consistent for i with P^0 and s^0 , for every $i \in N$. By consistency and optimality it follows that P^0 and $s_i^{(i)}$ satisfy optimality for i , and from consistency it follows that $s_i^{(i)} = s_i^*$, which concludes the proof.

The converse of Proposition ?? is essentially also true: if s^0 is a Nash equilibrium of (Γ, P^0) under a common prior P^0 , then there evidently exists an action-consistent equilibrium profile of conjectures C^* supporting s^0 , and such that players have the same first-order belief under C^* . Denote this common prior by P , and since P^0 and s^0 satisfy optimality for each $i \in N$, take $P = P^0$. For each $i \in N$ and $\ell \in \mathcal{L}^i$, let $C_i^\ell = (P^0, s^0)$, and the profile of conjectures thus defined C^* satisfies a common belief in optimality and consistency.²⁰

²⁰In his setting, Esponda (2011) gives a result analogous to Proposition ??, see Remark 1, p. 14.

3 A descriptive comparison with the common-prior benchmark.

The purpose of this section is to formulate precisely in what sense assuming the common-prior is restrictive as a positive description of equilibrium. If an outside observer, an analytical economist or econometrician, made the assumption of the common prior, when could that be a restrictive assumption, leading to mistaking some equilibrium outcomes (in the above sense) for non-equilibrium behavior?

There are many common environments where the common-prior assumption is not restrictive at all, for example, in private value environments, or in so-called generic environments (i.e., where each player's payoff attains a different value for every draw or types and action profile). In the former case, whenever a player optimizes against the actual play of her opponents under some belief, she is also optimizing under the true distribution of Nature's moves. In the latter case, action-consistency implies that a player's belief must equal the true distribution of Nature's moves. In both of these cases it is then true that an outcome mapping of an action-consistent Nash equilibrium is also an outcome mapping of a Nash equilibrium under the true common prior. In contrast, in our main example of Section ?? the common-prior assumption is restrictive in a much stronger sense: there does not exist any common prior such that the distribution over players' actions in a Nash equilibrium under that common prior would match the distribution over players' actions in the particular action-consistent Nash equilibrium. The present description of strategic interaction then provides a richer descriptive tool.

To formalize this, a perspective of an analytical economist is assumed – an *observer*, who observes the strategic interaction under idealized conditions. The set of assumptions imposed on the observer may be in principle different from the assumptions on

the players in the game (Γ, P^0) , as described in the previous sections. Nevertheless, a disinterested observer should be “properly informed” in the sense that he should not be able to observe more than what the players observe about each other.²¹ It is here assumed that the observer knows Γ , knows Γ is common knowledge among the players in N , and knows that their strategies are supportable in an action-consistent Nash equilibrium of Γ .²² The observer’s observation is ideal in that he has at his disposal a data set that can perfectly identify some aspects of the interaction. The following question can then be formulated:

In an action-consistent Nash equilibrium, can the observer explain the data under the assumption of a common prior?

In the most ideal case the observer knows the joint distribution over players’ types and actions. This can be thought of either as an idealized benchmark case, whereby the observer knows the *true* model of the world. The question is then whether the observed behavior can be explained as equilibrium interaction under the true distribution of Nature’s moves. The assumption of a common prior would then seem particularly innocuous, since players’ behavior would then truly be as if they had beliefs equal to P^0 .

Another possibility, which is congruent with the assumptions on players’ observations is that the observer only knows the distribution over players’ actions. Then

²¹In part I of his formulation of Bayesian games, Harsanyi (1967) discusses that the game should be considered from the perspective of a “properly-informed” observer, who does not know more than the players know. One could in particular imagine this to be a necessary condition for an observer with ulterior interests, such as a *mechanism designer*, in order to preserve trust and to avoid incentive issues – such considerations are at the heart of Holmstrom and Myerson (1983).

²²That is, the observer has been able to separately conduct for each player $i \in N$ revealed-preference experiments in order to determine i ’s cardinal utilities over all possible outcomes in Γ , and that these preferences have been made publicly known, i.e., common knowledge. Also note that all of this is under the assumption of the equilibrium steady state behavior. Even the behavior of players who might behave rationally once in the steady state would presumably display different patterns before they reached such a steady state. How to model such out-of equilibrium behavior is a question of learning or equilibrium convergence, which are not addressed here. See Fudenberg and Levine (1993b) for learning in recurrent games, and also Bernheim (1984) for a discussion of learning in the context of rationalizability.

the question is whether it is possible to find a common prior \tilde{P} , such that the observed distribution over players' actions is a NE under the common prior \tilde{P} , for any distribution over players' actions in an action-consistent Nash equilibrium. In this latter case, correct welfare evaluations are no longer possible, but the common-prior model still provides a positive explanation of players' behavior. These two notions of descriptive equivalence are formulated as properties of the game Γ , since the true distribution of Nature's draws P^0 may or may not be known to the observer.

DEFINITION 2. Let s^* be supportable in an action-consistent Nash equilibrium of (Γ, P^0) .

1. s^* satisfies *outcome-equivalence*, if s^* is a NE under the common prior P^0 in (Γ, P^0) .
2. s^* satisfies *action-equivalence*, if there exists a common prior \tilde{P} , and a strategy profile \tilde{s} , such that \tilde{s} is a NE under the common prior \tilde{P} , and,

$$G[s^*, \hat{P}]_A = G[s^*, \tilde{P}]_A.$$

Γ satisfies *outcome-equivalence* if every profile of strategies supportable as an action-consistent Nash equilibrium of (Γ, P^0) satisfies outcome-equivalence, for every P^0 . Γ satisfies *action-equivalence* if every profile of strategies supportable as an action-consistent Nash equilibrium of (Γ, P^0) satisfies action-equivalence, for every P^0 .

It is noted that,

$$\text{outcome-equivalence} \Rightarrow \text{action-equivalence},$$

while the reverse implication is not true. If action-equivalence is not satisfied for some Γ , then neither will be outcome-equivalence. Hence, action-equivalence serves

as a stronger test as to how restrictive the common-prior assumption is.

Action-equivalence seems a very natural observation criterion for the observer. In econometric applications, data on observations of players' actions seems natural, and available; of course, often only coarser data might be available, but even then the data on players' actions may represent an ideal benchmark. In certain theoretical applications too, in particular, in adverse-selection environments, the assumption is that the data on players' actions is available but not the data on their types. Whenever action-equivalence is used for the purpose of positive analysis, action-consistency is the notion of consistency congruent with Harsanyi's idea of a properly-informed observer. Effectively, any one of the players of the game itself, can be thought of as such a properly-informed observer.

As mentioned at the beginning of the section, examples of environments where outcome-equivalence obtains are games with private values and games with generic interdependence. Formally, a game satisfies *private values* if,

$$u_i(a, t_i, t_{-i}) = u_i(a, t_i, t'_{-i}), \forall t_{-i}, t'_{-i} \in T_{-i}, \forall t_i \in T_i, \forall a \in A.$$

A game satisfies *generic interdependence* if,

$$u_i(a, t_i, t_{-i}) \neq u_i(a, t_i, t'_{-i}), \forall t_{-i}, t'_{-i} \in T_{-i}, \forall t_i \in T_i, \forall a \in A.$$

When a game satisfies private values, a player need only to best-reply to the distribution over opponents' actions, so that her conjectures about marginal probability over opponents' types, and consequently about the joint distribution, are irrelevant. When a game satisfies generic interdependence, opponents' types nontrivially affect each player's payoffs. By consistency, the frequencies of her own payoffs are conjectured with correct probabilities, so that the player's conjectures regarding the joint

distribution of types must be correct, and outcome-equivalence obtains.²³ To be sure, in order for a player's beliefs to coincide with the true prior, a weaker condition than generic interdependence is needed: a player should be able to identify the true prior from the variation of her payoffs across types of other players, given her type. That is, for each draw of her type, a player should be able to identify conditional joint distribution from the variation in her own payoffs.

4 Main example

Next is the main example of this paper. It is given by a game $(\bar{\Gamma}, P^0)$. This game has the property that there exists an action-consistent Nash equilibrium, which is not a Nash equilibrium under any common prior. That is, $\bar{\Gamma}$ does not satisfy action equivalence.²⁴

Game $(\bar{\Gamma}, P^0)$ is given as follows. There are two players, $N = \{1, 2\}$, each player has two types, $T_i = \{t_i, t'_i\}$, $i \in N$, two actions, $A_1 = \{up, down\}$, $A_2 = \{L, R\}$, and the payoffs for each draw of types are specified in the following tables.

(t_1, t_2)	L	R
up	$-2, \underline{3}$	$\underline{1}, \underline{1}$
down	$0, 0$	$\underline{0}, 0$

(t_1, t'_2)	L	R
up	$-1, \underline{1}$	$\underline{1}, \underline{2}$
down	$-1, -1$	$\underline{3}, -3$

²³Note that in both of these environments, it is enough to specify hierarchies of assessments only up to the second order. This is then formally equivalent to the setting of Proposition 2, in Dekel et al. (2004) in the context of self confirming equilibrium. Their result is that, under the particular signaling structure equivalent to action-consistency, the self-confirming equilibrium is a NE under the true common prior, and the argument above follows their proof.

²⁴Note that this example is quite different from various examples in the literature on non-common priors in games. In that literature, either players' beliefs are not anchored at all, or if they are anchored, players do not make any inference from rationality of other players. The most similar to the present setup is that of Esponda (2011), but there only the actual draw of parameters by Nature is considered, and not the counterfactual draws of parameters. In the ensuing game as defined by Esponda, players' strategies are then probability distributions over actions, and not mappings from types into actions as in Harsanyi's definition of a game of incomplete information. That players' strategies are mappings is crucial for Harsanyi's approach and for the present analysis, as it is precisely here where the players' inability to identify the true distribution of Nature's draws and other's strategies comes from.

(t'_1, t_2)	L	R
up	$-1, \underline{0}$	$\underline{0}, \underline{1}$
down	$-1, -1$	$\underline{1}, -2$

(t'_1, t'_2)	L	R
up	$-3, \underline{3}$	$\underline{0}, \underline{2}$
down	$-2, -1$	$\underline{-1}, 0$

The objective probability distribution of Nature's moves P^0 is given by,

P^0	t_2	t'_2
t_1	$\frac{1}{4}$	$\frac{1}{4}$
t'_1	$\frac{1}{4}$	$\frac{1}{4}$

In game $(\bar{\Gamma}, P^0)$, the strategy profile $s^* = (up|t_1, up|t'_1; R|t_2, R|t'_2)$ is shown to be an *equilibrium strategy profile*, consistent with the true probability distribution of Nature's moves P^0 , when $\bar{\Gamma}$ is played by entirely rational and sophisticated players. However, this outcome is not an equilibrium outcome of $(\bar{\Gamma}, P^0)$ under *any common-prior* probability distribution of Nature's moves \underline{P} . Appropriate supporting players' beliefs and *conjectures* are constructed in order to show that this outcome is a *action-consistent Nash equilibrium* outcome of $(\bar{\Gamma}, P^0)$.

Consider the strategy profile given by $s^* = (up|t_1, up|t'_1; R|t_2, R|t'_2)$. First, there does not exist *any common prior* \underline{P} , such that s^* could be supportable in a NE of $(\bar{\Gamma}, P^0)$ under common prior \underline{P} . Suppose that there existed such a common prior \underline{P} . The following incentive constraints would then have to hold for the strategy profile s^* – in the above tables the payoffs which are relevant for these incentive constraints are underlined:

$$0 \geq 2\underline{P}(t_1, t'_2) - \underline{P}(t_1, t_2), \quad (3)$$

$$0 \geq \underline{P}(t'_1, t_2) - \underline{P}(t'_1, t'_2), \quad (4)$$

$$0 \geq 2\underline{P}(t_1, t_2) - \underline{P}(t'_1, t_2), \quad (5)$$

$$0 \geq \underline{P}(t'_1, t'_2) - \underline{P}(t_1, t'_2). \quad (6)$$

Adding these inequalities yields $3\underline{P}(t_1, t_2) \leq 0$, so that $\underline{P}(t_1, t_2) = 0$. By (??), it follows that $\underline{P}(t_1, t'_2) = 0$, implying by way of (??) that $\underline{P}(t'_1, t'_2) = 0$, which by way of (??) implies that $\underline{P}(t'_1, t_2) = 0$, so that $\underline{P} \equiv 0$, which is a contradiction.

Is s^* nonetheless supportable as an *equilibrium* of $(\bar{\Gamma}, P^0)$, when players have different beliefs?

So suppose that players' beliefs \bar{P}^1 and \bar{P}^2 regarding the joint probability distribution over types T are given by the following two matrices,

\bar{P}^1	t_2	t'_2
t_1	$\frac{4}{10}$	$\frac{1}{10}$
t'_1	$\frac{1}{10}$	$\frac{4}{10}$

\bar{P}^2	t_2	t'_2
t_1	$\frac{1}{10}$	$\frac{4}{10}$
t'_1	$\frac{4}{10}$	$\frac{1}{10}$

A superficial observation is that s^* and \bar{P}^1 satisfy incentive constraints (??) and (??), while s^* and \bar{P}^2 satisfy incentive constraints (??) and (??). Hence, if player 1 has prior \bar{P}^1 and player 2 has prior \bar{P}^2 , each of them is optimizing given their beliefs; One might be tempted to think of this as a definition of Nash equilibrium when there is no common prior. But *much more* must be true for P^0 and s^* , along with \bar{P}^1 , and \bar{P}^2 .

Under belief \bar{P}^i each player i *consistently* assesses the situation. For each of player i 's types, and given this type's action, i has rational expectations regarding the distribution over the other player's actions and his own payoffs. For example, when player 1 is of type t_1 and plays *up*, his belief $\bar{P}_{T_2|t_1, up}^1$ along with an assessment that player 2 plays action R implies that player 1 should obtain a payoff of 1 and observe player 2 playing action R with probability 1. This is indeed what happens under the actual distribution of Nature's moves P^0 , and the actual strategy profile s^* . Thus, both players have assessments that are consistent with what actually happens in terms of the conditional distribution over own payoffs, and the other

players' actions. Their beliefs are also consistent with the marginal distribution over own types, i.e., each player knows the true distribution of how likely he is to be of a particular type. Hence, action-consistency is satisfied.

Players' conjectures supporting the equilibrium outcome are constructed as follows. Player 1 may conjecture that: Nature's moves are given by \bar{P}^1 and that the two players are playing a strategy profile s^* ; That player 2's assessment of the situation is that Nature's moves are given by \bar{P}^2 and that the strategy profile is s^* ; That player 2's assessment about player 1's assessment is that Nature moves according to \bar{P}^1 , and the strategy profile is s^* , and so forth. Hence, player 1's conjecture is given by the hierarchy of assessments,

$$C_1^* = (P^{(1)*}, s^{(1)*}; P^{(1,2)*}, s^{(1,2)*}; P^{(1,2,1)*}, s^{(1,2,1)*}; \dots),$$

where $s^{(\ell)*} = s^*$, for every sequence ℓ of 1's and 2's, and $P^{(1)*} = \bar{P}^1$, $P^{(1,2)*} = \bar{P}^2$, $P^{(1,2,1)*} = \bar{P}^1$, and so forth. Similarly to C_1^* , player 2 can make the following conjecture,

$$C_2^* = (\bar{P}^2, s^*; \bar{P}^1, s^*; \bar{P}^2, s^*; \dots).$$

Conjectures C_1^* and C_2^* satisfy common belief in *action-consistency*. At any level, what player i presumes about the opponent is consistent with player i 's assessment at that level. For example, player 1's assessment regarding player 2's assessment about player 1 is consistent with player 1's assessment; it need not be consistent with P^0 , as player 1 does not know P^0 .

Conjectures C_1^* and C_2^* also allow both players to rationalize the situation, i.e., they satisfy common belief in optimality. Player 1 is able to rationalize 2's behavior, i.e., player 2 is optimizing under beliefs which 2 might have consistently with 1's observations. Player 1 is making an assessment under which player 2's strategy is *optimal from player 2's perspective*. Similarly, player 1's conjecture allows player 1

to rationalize the presumed assessment of player 2 regarding player 1, and so on. Players attribute optimal behavior to each other at any order of assessment.

Of course, each player could entertain a host of different conjectures. Some might satisfy consistency, some might rationalize the opponent's behavior, perhaps up to some order, and some might satisfy none of these properties. Players might also have conjectures which have the same properties as the ones above, but where players do not happen to attribute to each other their actual beliefs and strategies. Conjectures C_1^* and C_2^* constitute a particularly simple profile of conjectures satisfying all our requirements: no matter how far in her thinking a player went, she would always be able to presume that each of them is playing optimally and has assessments that are consistent with what is going on. Hence, no player would have any reason to change her conjecture, either on the basis of her own optimization, or on the basis of potential inconsistency with reality, when their observation of reality includes the opponent's actions in the game.

5 A finite-order example

In this section, game $\bar{\Gamma}$ is modified to illustrate the difference between the analysis based on a closed model (i.e., an infinite hierarchy of assessments) and an analysis based on a finite order of assessments. In our definition of action-consistent Nash equilibrium, hierarchies of assessments are infinite and belief in optimality and consistency must hold at any order. In the main example it is also necessary to specify such an infinite hierarchy of assessments in order for players to be able to rationalize each other's behavior: At any order, i 's belief only rationalizes her own strategy, but not the opponent's; for the opponent's strategy to be rationalized the next order of assessment must be specified. In the following example, game $\bar{\Gamma}$ from Section ?? is modified to game $\bar{\Gamma}'$, where a particular profile satisfies optimality and consistency up to any finite order (depending on P^0), but it does not satisfy optimality and con-

sistency in an infinite hierarchy of assessments. Hence, if the observer only observes the distribution over players' actions, then he can in game $\bar{\Gamma}'$ "explain" players' strategy profile (up, R) with any finite order of belief in consistency and optimality, but (up, R) is not an action-consistent equilibrium of $\bar{\Gamma}'$. This illustrates how from the observer's perspective a model with an infinite hierarchy of assessments is different from a model that is not closed.

Let the game $\tilde{\Gamma}'$ be given by the following payoff matrices,

(t_1, t_2)	L	R
up	$-2, \underline{\underline{3}}$	$\underline{1}, \underline{1}$
down	$0, 0$	$\underline{0}, 0$

(t'_1, t_2)	L	R
up	$-1, \underline{0}$	$\underline{0}, \underline{1}$
down	$-1, -1$	$\underline{1}, -2$

(t_1, t'_2)	L	R
up	$-1, \underline{\underline{3}}$	$\underline{\underline{0}}, \underline{\underline{0}}$
down	$-1, -1$	$\underline{\underline{2}}, -3$

(t'_1, t'_2)	L	R
up	$-3, \underline{1}$	$\underline{0}, \underline{0}$
down	$-2, -1$	$\underline{\underline{-1}}, 0$

Note that $\bar{\Gamma}'$ is almost the same as $\bar{\Gamma}$, except for the underlined payoffs in bold type. That is, in $\bar{\Gamma}'$, for the type draw (t_1, t'_2) , and player 2's action $a_2 = R$, player 1 obtains $\mathbf{0}$ when playing *up*, and $\mathbf{2}$ when playing *down*; recall that in $\bar{\Gamma}$, the corresponding payoffs were 1 and 3. Note that under this change, the incentive constraints for the strategy profile $(up|t_1, up|t'_1; R|t_2, R|t'_2)$ in $\bar{\Gamma}'$ remain identical to those in $\bar{\Gamma}$, i.e., they are given by (??)-(??). However, in $\bar{\Gamma}'$, when the profile is $(up|t_1, up|t'_1; R|t_2, R|t'_2)$, and player 1 is of type t_1 , her payoffs now vary with the types of player 2, so that in order for consistency to hold, 1 must correctly assess the probabilities of t_2 and t'_2 , conditional on t_1 .

Now suppose that the distribution of Nature's moves is given by P' ,

P'	t_2	t'_2
t_1	$\frac{7}{8}$	$\frac{1}{24}$
t'_1	$\frac{1}{24}$	$\frac{1}{24}$

Let $P^{(1)} = P'$, and let $P^{(2)}$ be given by,

$P^{(2)}$	t_2	t'_2
t_1	$\frac{7}{24}$	$\frac{1}{24}$
t'_1	$\frac{15}{24}$	$\frac{1}{24}$

Next, let $P^{(12)} = P^{(2)}$, and let $P^{(21)}$ be given by,

$P^{(21)}$	t_2	t'_2
t_1	$\frac{7}{24}$	$\frac{1}{24}$
t'_1	$\frac{8}{24}$	$\frac{8}{24}$

Finally, let $P^{(121)} = P^{(21)}$, and let $P^{(212)}$ be given by,

$P^{(212)}$	t_2	t'_2
t_1	$\frac{5}{24}$	$\frac{9}{48}$
t'_1	$\frac{10}{24}$	$\frac{9}{48}$

Up to this level, that is 3, these assessments satisfy all the consistency requirements, and also satisfy the incentive constraints, i.e., rationalizability. At the next order, it is still possible to define $P^{(1212)} = P^{(212)}$. However, it is impossible that $P^{(2121)}$ could satisfy the incentive constraints, as $2 \times \frac{9}{48} - \frac{5}{24} = \frac{4}{24} > 0$. It is evident that this example can be modified so that rationalizability and consistency are satisfied up to an arbitrarily high finite order, but can never be satisfied in an

infinite hierarchy of assessments.²⁵ This last example illustrates how common belief in action-consistency and optimality, is different from a belief in action-consistency and optimality up to any finite order.²⁶

²⁵More precisely, if under \hat{P}' the probability of (t_1, t_2) is $\frac{2^k-1}{2^k}$, while the other types have probabilities equal to $\frac{2^k-1}{3 \times 2^k}$, then it can be shown that consistency and rationalizability will hold up to order $k-1$ for both players, but will fail for at least one player at the order k .

²⁶Assessments satisfying optimality and action-consistency up to order 1 are equivalent to the notion of self-confirming equilibrium, when the signalling function is given by the other players' actions and own payoffs (see Fudenberg and Levine (1993a), and Dekel et. al (2004)). Thus, in a self-confirming equilibrium players' are level-1 sophisticated, so that they do not question other players' motives beyond that. For example, in the above game $\bar{\Gamma}'$, under P' , the profile (up, R) would constitute a self-confirming equilibrium for this feedback function, and the set of distribution functions P' for which that would be the case would be quite large.

References

- [1] AUMANN AND BRANDENBURGER [1995]: “Epistemic Conditions for Nash Equilibrium,” *Econometrica*, **63**, 1161-1180.
- [2] BATTIGALLI, P. [1987]: “Comportamento Razionale Ed Equilibrio Nei Giochi e Nelle Situazioni Sociali,” unpublished undergraduate dissertation, Universit Bocconi, Milano.
- [3] BATTIGALLI, P., A. DI TILLIO, E. GRILLO, AND A. PENTA [2009]: “Interactive Epistemology and Solution Concepts for Games with Asymmetric Information, working paper.
- [4] BATTIGALLI P. AND D. GUATOLI [1997]: “Conjectural Equilibria and Rationalizability in a Game with Incomplete Information,” in *Decisions, Games and Markets*, Kluwer Academic Publishers, Norwell, MA.
- [5] BATTIGALLI, P. AND M. SINISCALCHI [2003]: “Rationalization and Incomplete Information, *Advances in Theoretical Economics* **3**, Article 3.
- [6] BERNHEIM, D. B. [1984]: “Rationalizable Strategic Behavior,” *Econometrica*, **52**, 1007-1028.
- [7] BRANDENBURGER, A. AND E. DEKEL [1987]: “Rationalizability and Correlated Equilibrium, *Econometrica* **55**, 1391-1402.
- [8] COPIC, J. AND A. GALEOTTI [2007]: “Awareness Equilibrium,” mimeo.
- [9] DEKEL E., D. FUDENBERG, AND D. K. LEVINE [2004]: “Learning to Play Bayesian Games,” *Games and Economic Behavior*, **46**, 282-303.
- [10] ESPONDA I. [2011]: “Rationalizable Conjectural Equilibrium,” mimeo.

- [11] FUDENBERG D. AND Y. KAMADA [2011]: “Rationalizable Partition-Confirmed Equilibrium,” mimeo.
- [12] FUDENBERG D. AND D. K. LEVINE [1993a]: “Self-Confirming Equilibrium,” *Econometrica*, **61**, 523-545.
- [13] FUDENBERG D. AND D. K. LEVINE [1993b]: “Steady State Learning and Nash Equilibrium,” *Econometrica*, **61**, 547-573.
- [14] HARSANYI, JOHN C. [1967]: “Games with incomplete information played by ‘Bayesian’ players, Part I – III., Part I. The basic model,” *Management Sci.*, **14**, 159182.
- [15] HARSANYI, JOHN C. [1968a]: “Games with incomplete information played by ‘Bayesian’ players, Part I – III., Part II. Bayesian equilibrium points,” *Management Sci.*, **14**, 320334.
- [16] HARSANYI, JOHN C. [1968b]: “Games with incomplete information played by ‘Bayesian’ players, Part I – III., Part III. The basic probability distribution of the game,” *Management Sci.*, **14**, 486502.
- [17] HOLMSTROM, BENGT AND ROGER B. MYERSON [1983]: “Efficient and durable decision rules with incomplete information,” *Econometrica*, **51**, 1799 -1819.
- [18] JACKSON M. O. AND E. KALAI [1997]: “Social Learning in Recurring Games,” *Games and Economic Behavior*, **21**, 102-134.
- [19] MERTENS J.-F. AND S. ZAMIR [1985]: “Formulation of Bayesian analysis for games with incomplete information,” *Internat. J. of Game Theory*, **14**, 129.
- [20] MORRIS S. [1995]: “The Common-Prior Assumption in Economic Theory,” *Economics and Philosophy*, **11**, 227-253.

- [21] NASH, J. [1951]: “Non-Cooperative Games,” *Annals of Mathematics* **54**, 286-295.
- [22] PEARCE, D. [1984]: “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica* **52**, 1029-51.
- [23] RADNER, R. [1982]: “Equilibrium under uncertainty,” Chap. 20 in *Handbook of Mathematical Economics*, vol. II, edited by K. Arrow and M. D. Intrilligator. Amsterdam: North-Holland.
- [24] RUBINSTEIN A. AND A. WOLINSKY[1994]: “Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability,” *Games and Economic Behavior*, **6**, 299-311.
- [25] SIMON, LEO K. AND WILLIAM R. ZAME [1990]: “Discontinuous Games and Endogenous Sharing Rules,” *Econometrica*, **58**, 861-872.
- [26] VICKREY, W. [1961]: “Counterspeculation, Auctions, and Competitive Sealed Tenders, *Journal of Finance*, **16**, 8-37.