Documents de travail

«Distance in Beliefs and Individually-Consistent Sequential Equilibrium»

Auteurs

Gisèle Umbhauer, Arnaud Wolf

Document de Travail n° 2022 – 37

Décembre 2022

Bureau d'Économie Théorique et Appliquée BETA

www.beta-economics.fr

@beta_economics

Contact :
jaoulgrammare@beta-cnrs.unistra.fr

Distance in Beliefs and Individually-Consistent Sequential Equilibrium

Gisèle Umbhauer* Arnaud Wolff†

December 2022

Abstract. The concept of Individually-Consistent Sequential-Equilibrium broadens the concept of Sequential Equilibrium by allowing players to have different beliefs on potential deviations. This heterogeneity spontaneously gives rise to a notion of distance between beliefs. Yet, studying distance between beliefs in a strategic context reveals to be intricate. Announced beliefs may be different from revealed beliefs and the meaning of distance depends on the role assigned to beliefs. If out-of-equilibrium beliefs help getting a larger payoff at equilibrium, then we might need to reconsider the traditional definition of sequential rationality: more than just requiring that players behave optimally at every information set given their beliefs and the strategies played by others players, we might additionally require that there does not exist another perturbation scheme that is individually-consistent and which provides higher payoffs to the players.

Keywords: AGM-Consistency, Distance in Beliefs, Heterogeneous Beliefs, Individually-Consistent Sequential Equilibrium, Revealed Beliefs

JEL: C72

*Bureau d'Economie Théorique et Appliquée (BETA) - Université de Strasbourg - 61 Avenue de la Forêt Noire, 67000 Strasbourg - France. e-mail: umbhauer@unistra.fr.

†Bureau d'Economie Théorique et Appliquée (BETA) - Université de Strasbourg - 61 Avenue de la Forêt Noire, 67000 Strasbourg - France. e-mail: arnaudwolff@unistra.fr.

1. Introduction

In this paper, we extend the concept of Individually-Consistent Sequential Equilibrium (ICSE, Umbhauer and Wolff 2019), which builds on the Sequential Equilibrium (SE, Kreps and Wilson 1982), a solution concept commonly used to solve extensive-form games. The SE requires consistency of beliefs at all information sets, even at those that find themselves out of the equilibrium strategy path. This implies that players are required to share the same beliefs at out-of-equilibrium information sets, even about the numerical values of mathematical artifacts used to generate perturbations of strategy profiles, which are arbitrary by nature. Since there is no *a priori* basis for requiring players to agree on the probabilities of other players' possible mistakes (or deviations), the ICSE accepts different perturbation systems for different players.

This paper focuses on games with $n \geq 3$ players since these are the games in which the ICSE solution concept can differ from the SE. In particular, we focus on games in which some players might belong to a same social group or a same community. Therefore, although out-of-equilibrium beliefs are never directly confronted to reality (so that players can in some sense *agree to disagree*), players that belong to a same social group may feel ill at ease when adopting different beliefs. In fact, research in political science has shown that individuals are often motivated to shift their beliefs towards the ones associated with the social groups they belong to (Barber and Pope, 2019, Gould and Klor, 2019, Slothuus and Bisgaard, 2021). Therefore, rather than being completely arbitrary, beliefs at out-of-equilibrium information sets might be correlated among players. This leads us to develop a notion of distance between the beliefs of different players as well as the idea of maximally allowed heterogeneity between the players' beliefs.

The notion of distance between beliefs introduced in this paper can not be properly studied without further delving into the function of beliefs and their intrinsic link to actions. For instance, a player may publicly declare to hold some beliefs but his actions contradict the proclaimed beliefs. That is to say, the *revealed* beliefs of the player are different from the *announced* beliefs. The question then arises as to whether we should measure the distance between the player's announced beliefs or between their revealed beliefs. Furthermore, in a game as in real life, the purpose of out-of-equilibrium beliefs may be to help a player maximize his own payoffs. In this sense, beliefs become strategic and they may in some way belong to the strategy set of the players. For instance, it might be in Player 1 and Player 2's interests to have similar (or different) beliefs so as to incentivize Player 3 to adopt a strategy that maximizes their own payoffs. In this context, players build their beliefs with a strategic purpose and the distance between their beliefs becomes irrelevant. These considerations lead us to revisit the notion of sequential rationality in dynamic games of incomplete information. More than just requiring that players behave optimally at every information set given their beliefs and the strategies played by other players, we might additionally require that there does not exist another perturbation scheme that is individually-consistent and which provides higher payoffs to the players.

The paper is organized as follows. In Section 2, we describe the concept of ICSE and discuss the way it introduces heterogeneity in beliefs at out-of-equilibrium information sets. In Section 3, we compare our concept with other often-used solution concepts such as PBE or AGM-consistency. In Section 4, we turn to the notion of distance between beliefs. We present two ways of measuring distance. The first is an order relation on beliefs, while the second is an Euclidean notion of distance, in that we measure the minimal payoff perturbations necessary to ensure convergence in beliefs. Yet, the main difficulties remain elsewhere. In Section 5, we introduce the distinction between revealed and announced beliefs and discuss the strategic function of beliefs. Section 6 discusses the findings of this paper by revisiting the definition of sequential rationality. The last section concludes.

2. Individually-Consistent Sequential Equilibrium

In this paper, we consider finite extensive-form games and focus on games with $n \geq 3$ players. Let N represent the finite set of players (with typical element $n \in N$), X the set of non-terminal decision nodes (with typical element $x \in X$) and H the set of all possible information sets (with $h \in H$ a specific information set). Let $H_i \subseteq H$ denote the set of all possible information sets at which Player i might be called upon to play. We call $i(h)$ the player playing at h and for every $h \in H_i$, we note A_h the set of actions available to player i at information set h.

A *behavioral strategy* for player i, noted π_i, is a probability distribution over her possible actions at each of her information sets. That is, a behavioral strategy for player i is a member of $\times_{h \in H_i} \Delta(A_h)$. The set of behavioral-strategy profiles is therefore $\times_{i \in N} \times_{h \in H_i} \Delta(A_h)$, with typical element $\pi = (\pi_i)_{i \in N}$. A *system of beliefs* is a function $\mu : X \to [0, 1]$ such that $\forall h \in H, \Sigma_{x \in h} \mu(x) = 1$.

The Sequential Equilibrium (SE) requires *consistency* of beliefs at all information sets, even at those that find themselves out of the equilibrium strategy path. To generate beliefs that are consistent at every information set, Kreps and Wilson (1982) require that the belief vector μ be the limit of a sequence of belief vectors derived from Bayes' rule applied to a sequence of *fully mixed* strategy profiles (strategy profiles that put positive probability to every action in every information set). Let us denote by $\times_{h\in H}\Delta^0(A_h)$ the set of all fully mixed behavioral strategies. Formally, a pair (μ,π) is consistent if and only if there exists some sequence $(\hat{\mu}^k,\hat{\pi}^k)_{k=1}^{\infty}$ such that:

(i) $\hat{\pi}^k \in \times_{h\in H}\Delta^0(A_h), \forall k \in \{1,2,3,...\}$,

(ii) $\hat{\mu}^k_h(x) = \frac{P(x|\hat{\pi}^k)}{\sum\limits_{y\in h} P(y|\hat{\pi}^k)}, \forall h \in H, \forall x \in h, \forall k \in \{1,2,3,...\}$,[1]

(iii) $\pi_{i(h)}(a_h) = \lim\limits_{k\to\infty} \hat{\pi}^k(a_h), \forall i \in N, \forall h \in H, \forall a_h \in A_h$,

(iv) $\mu_h(x) = \lim\limits_{k\to\infty} \hat{\mu}^k_h(x), \forall h \in H, \forall x \in h$.

A SE is defined to be any pair (μ,π) that is consistent and sequentially rational (Kreps and Wilson 1982, p.872).

What is crucial in the concept of SE is that the players are required to implicitly agree on the value of the ϵ used to generate perturbations of the strategy profiles. That is, while the ϵ are arbitrary in nature (they only represent mathematical artifacts), players still need to share the same beliefs about their numerical values. In some way, it is as if an external player shakes the strategies for everybody. In Umbhauer and Wolff (2019), we argue that this requirement is too strong. Indeed, we argue that there is no *a priori* basis for requiring that players agree on the probabilities of other players' possible mistakes (or deviations).

Formally, what distinguishes our Individually-Consistent Sequential Equilibrium (ICSE) concept from the SE is that we do not require the existence of only one sequence of perturbed strategy profiles on which all players need to agree but allow for different perturbation systems for different players. In other words, each player j introduces his own perturbations on the actions at each information set $h \in H$. So, $\hat{\pi}^k_{j,i(h)}(a_h)$ is the value player j assigns to the probability with which player $i(h)$ plays a_h at his information set h, while $\hat{\pi}^k_j$ is player j's profile of perturbed strategies in the whole game. Of course, for consistency, we require that $\pi_{i(h)}(a_h) = \lim\limits_{k\to\infty}\hat{\pi}^k_{j,i(h)}(a_h)$ for all $j \in N$, so that each player j's perturbed strategy profile has to fit with the played actions in the game.

[1]With $P(x|\cdot)$ being computed using Bayes' rule.

Therefore, a pair (μ, π) is *individually-consistent* if and only if there exist some sequences $(\hat{\mu}_j^k, \hat{\pi}_j^k)_{k=1}^{\infty}$, for all $j \in N$, such that:

(i) $\hat{\pi}_{j,i(h)}^k \in \times_{h \in H} \Delta^0(A_h), \forall k \in \{1, 2, 3, ...\}, \forall j \in N$,

(ii) $\hat{\mu}_{i(h)}^k(x) = \frac{P(x|\hat{\pi}_{i(h)}^k)}{\sum\limits_{y \in h} P(y|\hat{\pi}_{i(h)}^k)}, \forall h \in H, \forall x \in h, \forall k \in \{1, 2, 3, ...\}$,

(iii) $\pi_{i(h)}(a_h) = \lim\limits_{k \to \infty} \hat{\pi}_{j,i(h)}^k(a_h), \forall h \in H, \forall a_h \in A_h, \forall j \in N$,

(iv) $\mu_{i(h)}(x) = \lim\limits_{k \to \infty} \hat{\mu}_{i(h)}^k(x), \forall h \in H, \forall x \in h$.

An *Individually-Consistent Sequential Equilibrium* (ICSE) is any pair (μ, π) that is both individually-consistent and sequentially rational.

Let us illustrate the consequences of such a concept. In the game in Figure 1, there does not exist any SE leading Player 1 to play C_1 (see Appendix 1), so the players can not reach the Pareto optimal payoffs (5.99, 10, 10). As a matter of fact, to sustain B_2, Player 2 has to believe that Player 1 trembles toward B_1 at least 4 times more often than toward A_1 ($\mu(x_2) \leq \frac{1}{5}$), whereas to be willing to play B_3, Player 3 has to believe that Player 1 trembles toward B_1 at most 3 times more often than toward A_2 ($\mu(y_2) \geq \frac{1}{4}$, hence $\mu(x_2) \geq \frac{1}{4}$). This is not possible in a SE, in that all the players shake the strategies in the same way. Yet this becomes possible with the ICSE.

What is new, in comparison with the SE, is the fact that players can have different beliefs at the same out-of-equilibrium information set. So, in the above example, we can set: $\mu_2(x_2) = 0.1$, $\mu_2(x_3) = 0.9$ for Player 2, and $\mu_3(y_1) = 0$, $\mu_3(y_2) = 0.3$, $\mu_3(y_3) = 0$ and $\mu_3(y_4) = 0.7$ for Player 3. This implicitly means that Player 3 assigns the belief 0.3 to x_2 and the belief 0.7 to x_3, given that Bayes' Rule requires that Player 3's beliefs are consistent: $\mu(y_2) = \mu(x_2)$ and $\mu(y_4) = \mu(x_3)$. This is due to the fact that each player shakes the strategies in the way he wants. So, for example, Player 2 may have in mind the perturbed strategy profile $\{(1 - \epsilon^k - 9\epsilon^k)C_1 + 9\epsilon^k B_1 + \epsilon^k A_1, (1 - \epsilon^k)B_2 + \epsilon^k A_2, (1 - \epsilon^k)B_3 + \epsilon^k A_3\}$, while Player 3 may have in mind $\{(1 - 3\epsilon^k - 7\epsilon^k)C_1 + 7\epsilon^k B_1 + 3\epsilon^k A_1, (1 - \epsilon^k)B_2 + \epsilon^k A_2, (1 - \epsilon^k)B_3 + \epsilon^k A_3\}$.[2] Heterogeneous beliefs at out-of-equilibrium information sets can therefore sustain the Pareto optimal payoffs at equilibrium in this strategic context.

[2] Player 1's perturbations have no impact on the game, so we can suppose that he has the same profile of perturbed strategies as Player 2 for example.

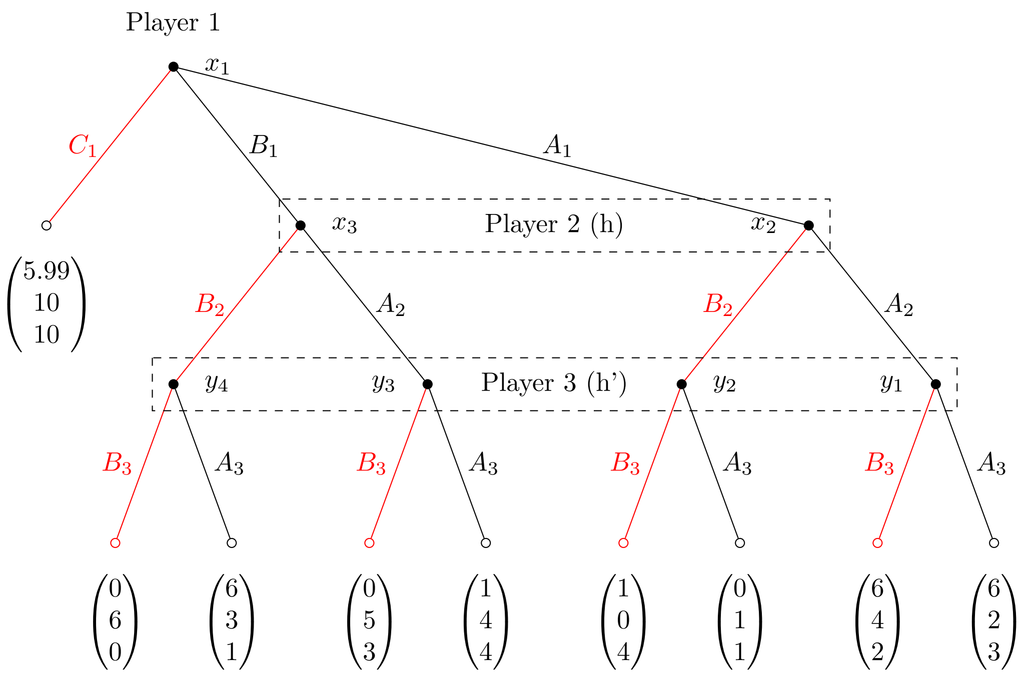

Figure 1: An example of the distinction between ICSE and SE.

3. Connections between ICSE and other solution concepts

3.1. Links between ICSE, SE, SPNE, PBE and SCE

Few solution concepts support the idea that people may share different beliefs at out-of-equilibrium information sets.[3] In fact, researchers tend to require that "players can not agree to disagree", the logical result stemming from Aumann (1976)'s paper. Yet, in our context, there is no true state to discover since the beliefs are about deviations that will never occur at equilibrium. So Player 2 (respectively Player 3) can durably think that, if they should face a deviation, then surely Player 1 played A_1 with a probability lower than 1/5 (respectively with a probability larger than 1/4). Nothing will contradict their beliefs given that Player 1 never deviates.

In the following Proposition, we enumerate the existing links between our ICSE solution concept and other well-known and often-used solution concepts, such as Subgame-Perfect Nash Equilibrium (SPNE), Sequential Equilibrium (SE), Perfect Bayesian Equilibrium (PBE) and Self-Confirming Equilibrium (SCE).

[3]We thank Giacomo Bonanno for informing us that Greenberg et al. (2009) developed a similar idea in their MACA concept.

Proposition 3.1. *(i) The set of ICSE is included in the set of Subgame-Perfect Nash Equilibria (SPNE).*

(ii) By construction, the set of SE is included in the set of ICSE, so the existence of an ICSE in a finite extensive-form game follows from the existence of a SE in a finite extensive-form game.

(iii) The set of ICSE is equal to the set of SE in a two-player game.

(iv) There is no inclusion relation between the set of Perfect Bayesian Equilibria (Fudenberg and Tirole, 1991) and the set of ICSE.

(v) The set of ICSE is included in the set of Self-Confirming Equilibria (Fudenberg and Levine, 1993).

Proof. In Appendix 2. ■

3.2. Links between ICSE and AGM-consistency

The ICSE also shares some links with the concept of AGM-consistency (Bonanno, 2013, 2016). AGM-consistency introduces a plausibility order on stories of actions and belief revision is based on this plausibility. This plausibility concept grants a large degree of freedom to the way beliefs are computed after deviations; this liberty differs from heterogeneous perturbation systems but it shares a partial link with the ICSE. The following Proposition describes these links.

Proposition 3.2. *i) The set of ICSE beliefs is almost included in the set of AGM-consistent beliefs.*

ii) There is no inclusion relationship between the set of ICSE beliefs and Bonanno's Perfect Bayesian Equilibrium beliefs.

Proof. In Appendix 2. ■

We come back to Bonanno's concept when studying the notion of distance, so we illustrate this concept on the game in Figure 1. Consider the ICSE with $\mu_2(x_2) = 0.1$, $\mu_2(x_3) = 0.9$, $\mu_3(y_1) = 0$, $\mu_3(y_2) = 0.3$, $\mu_3(y_3) = 0$ and $\mu_3(y_4) = 0.7$. These probabilities are compatible with AGM-consistency since they give positive weight to the actions (stories) A_1 and B_1 and to the stories A_1B_2 and B_1B_2. Therefore, they respect the plausibility-preserving action B_2. AGM-consistency is a qualitative notion, so the values of the beliefs are not important. What matters is that if the support of the beliefs are the stories A_1 and B_1, then the support of the stories reaching h' are the stories A_1B_2 and B_1B_2. So, every ICSE, which by definition lead to $\mu_2(x_2) \leq \frac{1}{5}$, $\mu_2(x_3) = 1 - \mu_2(x_2)$, and $\mu_3(y_1) = 0$, $\mu_3(y_2) \geq \frac{1}{4}$, $\mu_3(y_3) = 0$, $\mu_3(y_4) = 1 - \mu_3(y_2)$ respect AGM-consistency, except for the assessment that puts a 0 on $\mu_(x_2)$ or a 1 on $\mu_(y_2)$.

As a matter of fact, let us consider the "extreme" ICSE with $\mu_2(x_2) = 0$, $\mu_2(x_3) = 1$, $\mu_3(y_1) = 0$, $\mu_3(y_2) = 1$, $\mu_3(y_3) = 0$ and $\mu_3(y_4) = 0$. According to AGM-consistency, plausible histories can not sustain these beliefs, given that if $\mu(B_1)$ (the probability assigned to story B_1) is equal to 1 (because $\mu(x_2) = 0$ and $\mu(x_3) = 1$), then $\mu(B_1B_2)$ (the probability assigned to story B_1B_2) is also 1, since B_2 is played with probability 1 (it is the plausibility-preserving action); so the story B_1B_2 is as plausible as the story B_1. Given that A_1 is a less plausible story (in fact $\mu(A_1) = 0$) and given that $\mu(A_1) = \mu(A_1B_2)$, we get $\mu(A_1B_2) = \mu(y_2) < \mu(B_1B_2) = \mu(y_4)$, so $\mu(y_2)$ can not be equal to 1. With AGM-consistency, all happens as if an external observer deals with the possible beliefs of every player, upholding the planned equilibrium actions (as in the SE and in the ICSE) but possibly changing his view on an earlier out-of-equilibrium way of playing each time he faces a new deviation. So, if by observing that Player 1 does not play C_1, he becomes convinced that he plays B_1 ($\mu(x_3) = 1$), then, given that Player 2 plays B_2 at equilibrium, he necessarily assigns belief 1 to y_4.

We now consider Bonnano's PBE concept. The above ICSE, with $\mu(x_2) = 0.1$, $\mu(x_3) = 0.9$, $\mu(y_1) = 0$, $\mu(y_2) = 0.3$, $\mu(y_3) = 0$, and $\mu(y_4) = 0.7$ is not a PBE (Bonnano's version) in that, via Bayes' Rule, $\mu(x_2) = 0.1$ and $\mu(x_3) = 0.9$ lead to $\mu(y_1) = 0$, $\mu(y_2) = 0.1$, $\mu(y_3) = 0$ and $\mu(y_4) = 0.9$.

We finally show that, conversely, many AGM-consistent stories, and even Bonanno's PBE consistent stories, are not compatible with the concept of ICSE. So consider the game in Figure 2. Assume that the planned actions are the bold lines (in red) and that the beliefs (in blue) are given by $\mu(x_2) \geq 0.7$, $\mu(x_3) = 1 - \mu(x_2)$, $\mu(y_1) = \mu(x_2)$, $\mu(y_2) = 1 - \mu(x_2)$, $\mu(y_3) = \mu(y_4) = 0.5$.

These beliefs are AGM-consistent and they check Bonanno's PBE consistency. This is due to the fact that Bayes' Rule applies when switching from h to h' but it does not apply when switching from h to h'', since B_2 is not a plausibility-preserving action (by contrast to A_2). An external observer, when observing the unexpected action B_2, may completely reconsider the stories of the game. At h, he believes that Player 1 more often deviates to A_1 than to B_1, but after observing the new deviation B_2, he changes his mind and thinks that Player 1 deviates to B_1 as often as to A_1. This is not possible in an ICSE because the same player, Player 3, plays at h' and h''. Regardless of Player 3's perturbations on Player 1 and Player 2's actions, we necessarily

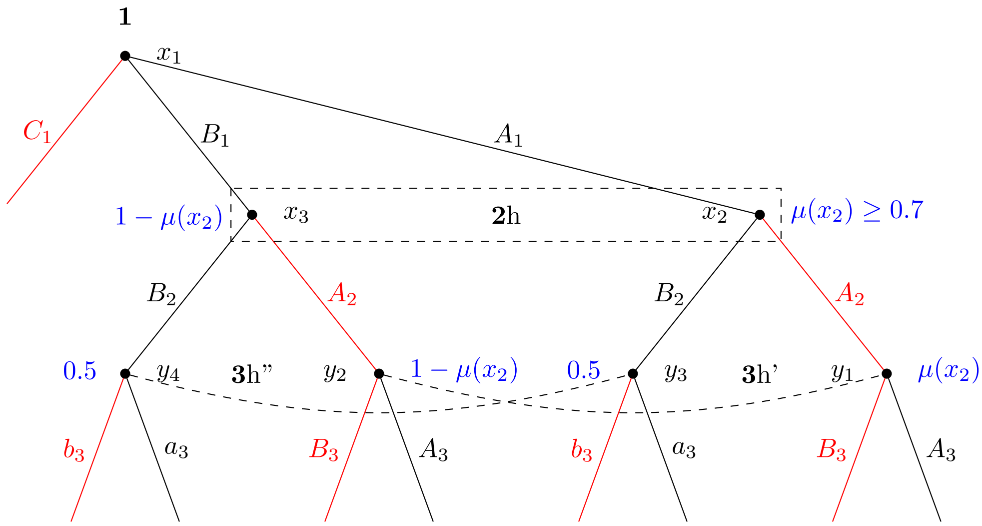

Figure 2: An example illustrating the difference between ICSE and AGM-consistency.

have $\mu_3(y_1) = \mu_3(y_3)$ and $\mu_3(y_2) = \mu_3(y_4)$, because Player 2 plays A_2 with the same probability at x_2 and x_3 and B_2 with the same probability at x_2 and x_3. Therefore, AGM-consistency allows an external player to have different *evolving* beliefs at a same information set (before Player 2's deviation, the external player sets $\mu(x_2) \geq 0.7$, but after Player 2's deviation he sets $\mu(x_2) = 0.5$), whereas the ICSE does not allow evolving beliefs at a same information set. Rather, it only allows different beliefs among the players (so $\mu_3(y_1) = \mu_3(y_3) = \mu(x_2)$ because these three probabilities express the way Player 3 evaluates the deviation from Player 1 towards A_1 and B_1 (before and after Player 2's choice of action), but Player 3's way of evaluating Player 1's deviations may be different from Player 2's way of evaluating these deviations, that is to say $\mu_2(x_2)$ can be different from Player 3's beliefs on x_2).

4. A physical distance between beliefs

While we defend the point of view that there is no logical reason that constrains people to have the same beliefs with respect to out-of-equilibrium actions, social groups often (implicitly, if not explicitly) require their members to have rather similar beliefs (Barber and Pope, 2019, Gould and Klor, 2019, Slothuus and Bisgaard, 2021), so the pressure to modify one's beliefs is increasing in the difference between the player's beliefs and the beliefs of the group they belong to. Therefore, if all the players in a game belong to a same community, it makes sense for them to seek to reduce the distance between beliefs.

To approach the distance between beliefs, we start with a first observation. In a game, very often, the equilibrium payoffs are sustained by sets of beliefs. For example, the ICSE equilibrium payoffs (5.99, 10, 10) in the game in Figure 1 are sustained by Player 2's beliefs $\mu_2(x_2) \leq \frac{1}{5}$ and Player 3's beliefs $\mu_3(y_2) \geq \frac{1}{4}$. So Player 2 has to assign a probability lower than $\frac{1}{5}$ to Player 1's deviating action A_1 whereas Player 3 has to assign a probability larger than $\frac{1}{4}$ to this deviation, a fact we reproduce in Figure 3, in which we highlight Player 2 and Player 3's sustaining beliefs (SB) on the action A_1.

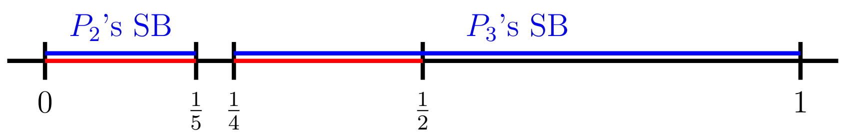

Figure 3: Player 2 and Player 3's sustaining beliefs.

We are concerned with the closest possible beliefs sustaining an ICSE, here $\frac{1}{5}$ and $\frac{1}{4}$. Figure 3 illustrates three facts:

(i) First, if Player 2's and Player 3's sustaining beliefs have a non-empty intersection, then common beliefs can sustain the ICSE payoffs and there exists a sequential equilibrium with these payoffs. So the minimum distance would reduce to 0.

(ii) Second, Player 2's and Player 3's sustaining beliefs have an empty intersection, but both sets are close ($\frac{1}{4}$-$\frac{1}{5}$ is small in comparison to 1), so that few changes in the game may lead both sets to have a non-empty intersection.

(iii) Third, we observe that both players can assign a probability lower than $\frac{1}{2}$ to A_1 to sustain the ICSE outcome (red probabilities for Player 2 and Player 3).

4.1. Ordered ICSE

Our first way to consider distance starts with the third observation. Imagine that Player 2 and Player 3 meet and discuss together: they can easily agree on the fact that both think that Player 1 deviates more often to B_1 than to A_1 (red probabilities lower than $\frac{1}{2}$). Player 2 is sure that Player 1 deviates at least 4 times more often to B_1 than to A_1. Player 3 can agree that Player 1 deviates more often to B_1 than to A_1 but at most 3 times more often. So there is a possible consensus between Player 2 and Player

3 despite the fact that they can not have the same beliefs. This consensus is on the way they order the deviations: both players believe that Player 1 more often deviates to B_1 than to A_1, yet only the "intensity" of this deviation is not shared among them.

It derives from this observation that a soft notion of proximity between beliefs consists in requiring that all the players order the perturbed strategies at each information set in the same way.

Definition 1. *An ordered ICSE is an ICSE that checks the additional condition:*

v) $\forall h \in H, \forall a, a' \in A_h, \forall j, j' \in N, \hat{\pi}^k_{j,i(h)}(a) \geq \hat{\pi}^k_{j,i(h)}(a') \Rightarrow \hat{\pi}^k_{j',i(h)}(a) \geq \hat{\pi}^k_{j',i(h)}(a')$.

In the game in Figure 1, it is easy to find an ordered ICSE that sustains the equilibrium actions (C_1, B_2, B_3) and that checks Definition 1. As a matter of example, with Player 2's distribution $(\epsilon, 4\epsilon, 1-5\epsilon)$ on the actions A_1, B_1 and C_1 and with Players 3's distribution $(\epsilon, 3\epsilon, 1-4\epsilon)$ on the actions A_1, B_1 and C_1, beliefs converge to sustaining beliefs and $\epsilon = \hat{\pi}^k_{2,1(x_1)}(A_1) < 4\epsilon = \hat{\pi}^k_{2,1(x_1)}(B_1) < 1-5\epsilon = \hat{\pi}^k_{2,1(x_1)}(C_1)$ and $\epsilon = \hat{\pi}^k_{3,1(x_1)}(A_1) < 3\epsilon = \hat{\pi}^k_{3,1(x_1)}(B_1) < 1-4\epsilon = \hat{\pi}^k_{3,1(x_1)}(C_1)$, for ϵ close to 0_+.

Yet, not every ICSE equilibrium actions can be sustained by probabilities that check Definition 1. For example, if Player 3's payoff 4 after $A_1B_2B_3$ is replaced by the payoff 1.9, then (C_1, B_2, B_3) will still be an ICSE, but no ICSE checks the condition in Definition 1. This is due to the fact that we necessarily have $\hat{\pi}^k_{2,1(x_1)}(A_1) < \hat{\pi}^k_{2,1(x_1)}(B_1)$ since we need $\mu_2(x_2) \leq \frac{1}{5} < \mu_2(x_3)$, and $\hat{\pi}^k_{3,1(x_1)}(A_1) > \hat{\pi}^k_{3,1(x_1)}(B_1)$ since we need $\mu_3(y_2) \geq \frac{1}{1.9} > \frac{1}{2} > \mu_3(y_4)$.

4.2. How to make ICSE beliefs SE-consistent?

Our second way to consider distance consists in exploiting the small size of the interval between Player 2's sustaining beliefs and Player 3's sustaining beliefs (second observation). Clearly, with respect to the game in Figure 1, the ICSE payoffs could become SE payoffs (and so the minimum distance between beliefs could collapse) by changing the game in a very smooth way. By replacing the payoff 1 after $B_1B_2A_3$ by 0.87 and the payoff 6 after $B_1B_2B_3$ by 6.17, it is possible to build an ICSE that is also a SE. We get $\mu_2(x_2) = \frac{4.5}{20}$, $\mu_2(x_3) = \frac{15.5}{20}$, and $\mu_3(y_1) = 0$, $\mu_3(y_2) = \frac{4.5}{20}$, $\mu_3(y_3) = 0$, $\mu_3(y_4) = \frac{15.5}{20}$, that is to say Player 2 and Player 3 have the same beliefs on Player 1's deviations.

It follows from this observation that another way to study the proximity between beliefs consists in looking at how much we need to shake the payoffs in order to get an ICSE that is also a SE; that is, how much we should shake payoffs to get beliefs that are consistent in a SE way. In other terms, after observing that it is not possible to get a SE with the equilibrium actions of a given ICSE, we can look if small changes in payoffs can allow us to get a SE with the ICSE payoffs.

Definition 2. *The ISCE beliefs are close if they can become SE-compatible with very small changes in payoffs. In that sense, the distance in beliefs becomes the distance in payoffs required to change the ICSE equilibrium payoffs into SE equilibrium payoffs.*

The steps are the following ones. We start with ICSE equilibrium behavioral strategies that can not be part of a SE. Then we introduce variables that express changes in payoffs and we minimize the changes in payoffs under the constraint that the ICSE actions and the associated beliefs become a SE.

Let us illustrate the procedure for the game in Figure 1, which is represented again in Figure 4. A first observation is that it is always possible to change the payoffs in order to get a SE with the ICSE played actions. The optimization program makes sense only if it leads to small payoff changes. In that case, we can say that the ICSE beliefs are not much distant one from another. If so, players will not feel under pressure to change them, namely because in real life there is always some incomplete information on the exact payoffs, so the smoothly changed payoffs (needed to share the same beliefs) belong to the set of possible payoffs.

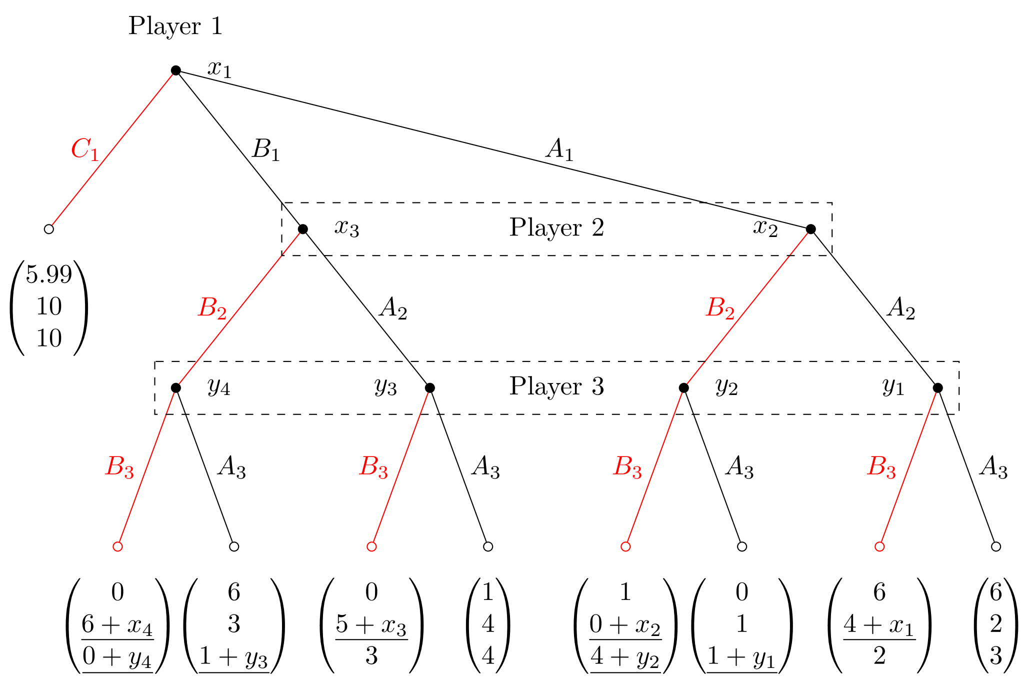

Figure 4: Absolute changes in payoffs that can make ICSE beliefs SE-compatible.

A second observation is that there are only a limited number of changes to introduce in that many payoffs have no role to play in the studied equilibrium. In our game, the payoffs to work on are the ones underlined. What matters is that, for (C_1, B_2, B_3) to become a SE, the SE consistency requires that $\mu(x_2) = \mu(y_2) = \mu$ and $\mu(x_3) = \mu(y_4) = 1 - \mu$. So the program we solve is Program 1:

$$\begin{aligned}
\min_{x_1,x_2,x_3,x_4,y_1,y_2,y_3,y_4,\mu} \quad & {x_1}^2 + {x_2}^2 + {x_3}^2 + {x_4}^2 + {y_1}^2 + {y_2}^2 + {y_3}^2 + {y_4}^2 \\
\text{s.t.} \quad & (0 + x_2)\mu + (6 + x_4)(1 - \mu) \geq (4 + x_1)\mu + (5 + x_3)(1 - \mu) \quad \mathbf{(1)} \\
& (4 + y_2)\mu + (0 + y_4)(1_\mu) \geq (1 + y_1)\mu + (1 + y_3)(1 - \mu) \quad \mathbf{(2)} \\
& \mu \geq 0 \quad \mathbf{(3)} \\
& 1 - \mu \geq 0 \quad \mathbf{(4)}
\end{aligned}$$

This objective function is one among the many possible ways to measure the changes in payoffs, surely the easiest one. We will propose later a more proportional one. Equations (1) and (2) are the necessary equations ensuring sequential rationality and SE consistency. Equation (1) ensures sequential rationality for Player 2, Equation (2) ensures sequential rationality for Player 3 and sequential rationality for Player 1 is ensured in that nothing has changed for himself with respect to Figure 1. What matters is that conditions (1) and (2) also ensure the SE consistency, which requires that Player 2 and Player 3 put the same belief μ on x_2 (and therefore on y_2), the same belief $1 - \mu$ on x_3 (and therefore on y_4) and a null belief on y_1 and y_3.

Given that the objective function goes to $+\infty$ when $\|x\|$ and/or $\|y\|$ goes to $+\infty$, and given that the admissible set is closed and that μ is limited by 0 and 1, it is easy to adapt Weierstrass' corollary to ensure that Program 1 has a global minimum. The only solution (see Appendix 3) is:

$$\begin{gathered}
x_2 = -x_1 \simeq 0.0158 \\
x_4 = -x_3 \simeq 0.0564 \\
y_2 = -y_1 \simeq 0.0206 \\
y_4 = -y_3 \simeq 0.0736 \\
\mu \simeq 0.219 \\
{x_1}^2 + {x_2}^2 + {x_3}^2 + {x_4}^2 + {y_1}^2 + {y_2}^2 + {y_3}^2 + {y_4}^2 = 0.0185
\end{gathered}$$

The necessary errors are quite small, since the largest one does not exceed 0.074, which is quite small given that we work with integers ranging from 0 to 6. In other terms, we can say that our ICSE payoffs are easily SE-compatible (because the needed payoffs changes are very small). Observe that the SE belief $\mu(x_2)$ becomes 0.219, which is between $\frac{1}{5}$ and $\frac{1}{4}$.

If we switch to a more proportional way to see payoff adjustments, we can choose to switch to Figure 5 and to the maximization Program 2:

$$\begin{aligned}
\min_{x_1,x_2,x_3,y_1,y_2,y_3,\mu} \quad & (\frac{x_1}{4})^2 + (\frac{x_2}{5})^2 + (\frac{x_3}{6})^2 + {y_1}^2 + (\frac{y_2}{4})^2 + {y_3}^2 \\
\text{s.t.} \quad & (6+x_3)(1-\mu) \geq (4+x_1)\mu + (5+x_2)(1-\mu) \quad \mathbf{(1)} \\
& (4+y_2)\mu \geq (1+y_1)\mu + (1+y_3)(1-\mu) \quad \mathbf{(2)} \\
& \mu \geq 0 \quad \mathbf{(3)} \\
& 1-\mu \geq 0 \quad \mathbf{(4)}
\end{aligned}$$

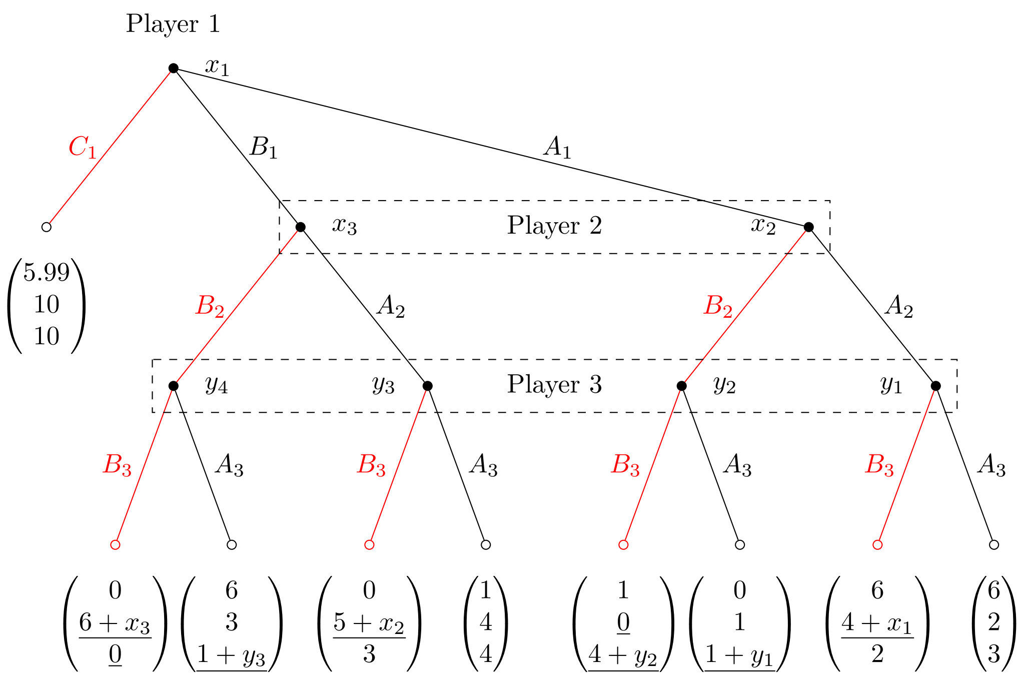

Figure 5: Proportional changes in payoffs that can make ICSE beliefs SE-compatible.

The only solution (see Appendix 3) is: $x_1 = -0.0258$, $x_2 = -0.1232$, $x_3 = 0.1774$, $y_1 = -0.0021$, $y_2 = 0.0338$, $y_3 = -0.0065$ and $\mu = 0.2466$.

We can observe that $(\frac{x_1}{4})^2 + (\frac{x_2}{5})^2 + (\frac{x_3}{6})^2 + {y_1}^2 + (\frac{y_2}{4})^2 + {y_3}^2 = 0.0016$, which is again quite small, and that no term (perturbation/payoff) is larger than 0.00645. The SE belief $\mu(x_1)$ now becomes 0.247, which is again between $\frac{1}{5}$ and $\frac{1}{4}$. Several remarks have to be made.

This concept of distance is rather simple to employ given that the program is easy to write (sequential rationality and the SE consistency give the set of constraints and the objective is a function increasing in the introduced payoff perturbations). But the interpretation of the result is necessarily a little subjective. For example, in Program 2, what should we require in order to say that beliefs are close? What is the threshold $|\frac{dx}{x}|$ we should accept (where dx is the variation of payoff and x the payoff)? We can of course impose the constraint $|\frac{dx}{x}| \leq 0.1$ in order to prevent too strong payoff changes, but this gives us no way to appreciate the optimal value of the objective function. Also, should we introduce a fixed threshold on the objective function and/or the ratios $|\frac{dx}{x}|$? Or should the thresholds depend on the payoffs that divergent beliefs allow to get at equilibrium? We think that, when opting for a proportional approach (Program 2), rather than asking for $|\frac{dx}{x}| \leq 0.1$ or another small value, we should ask for a threshold whose value rises with the benefit linked to the ICSE. This way of doing is motivated by the following fact: when a player earns a large payoff, he is less induced to change things (e.g., his beliefs) and he is less induced to ask that other persons change their beliefs.

4.3. Distance in ICSE beliefs and AGM-consistent beliefs

Another remark, which brings us back to Bonanno (2013, 2016)'s concept of AGM-consistency, is that this notion of distance does not take into account the number of deviations required to reach an information set. Let us consider the game in Figure 6. Suppose that it is possible to build an ICSE (C_1, A_2, B_3) with beliefs checking $\mu_2(x_2) \geq 0.6$ and $\mu_3(y_1) \leq 0.3$. Of course, SE consistency requires that $\mu_2(x_2) = \mu_3(y_1)$, so that we potentially have to strongly shake the payoffs in order to make the ICSE payoffs SE-compatible. In other terms, if we measure the distance as in Program 1 or Program 2, we will surely conclude that Player 2 and Player 3's beliefs are distant from one another. But this conclusion does not take into account a strong difference between h and h': h needs one deviation to be reached (Player 1's deviation) whereas h' needs two deviations to be reached (Player 1's deviation and Player 2's deviation).

This fact explains that Player 2 and Player 3's beliefs are AGM-consistent. As a matter of fact, the story A_1 is as plausible as the story A_1A_2 (because A_2 is the plausibility-preserving action), the story B_1 is as plausible as the story B_1A_2 for the same reason, and so an external observer can judge that the stories A_1 and A_1A_2 are more plausible than the stories B_1 and B_1A_2. Yet, he may also judge that the story B_1B_2 is more plausible than the story A_1A_2 because B_2 is a new deviation that totally

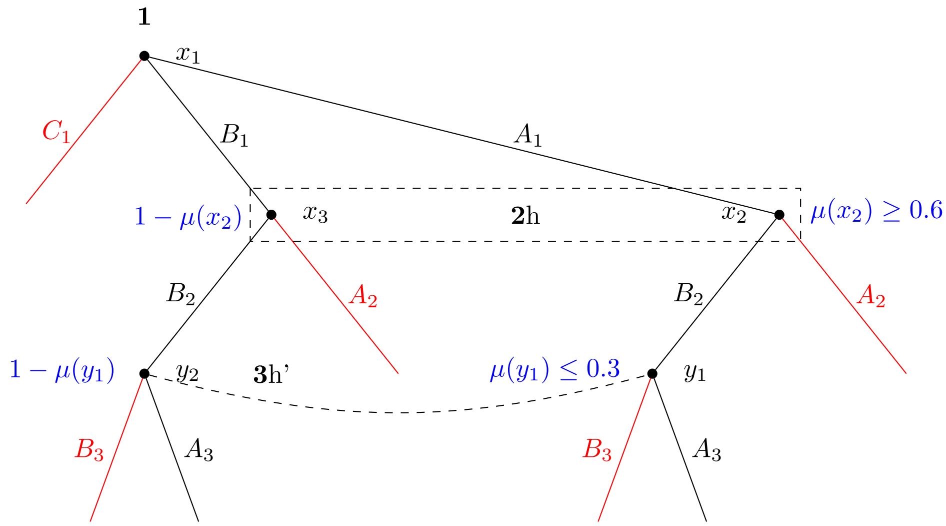

Figure 6: Revisiting the links between ICSE and AGM-consistency.

shakes his understanding of the game: at h, after Player 1's deviation, he thinks that Player 1 probably deviated toward A_1, so that Player 2 should play A_2, but after Player 2's unexpected deviation toward B_2 (at h'), he changes his mind and finally thinks that Player 1 probably deviated to B_1.

By contrast to Bonanno, we do not work with an external observer but only with the players themselves. This amounts, in Figure 6, in translating Bonanno's switch in beliefs into the following question: can we reasonably require that somebody who observes more deviations needs to share the same beliefs than somebody who observes less deviations? Player 2 observes only one deviation (Player 1's deviation) whereas Player 3 observes two deviations (Player 1's and Player 2's deviations). Facing more deviations may introduce more doubts and therefore allow for different beliefs. To say it differently, the distance in beliefs might also depend on the distance (in the number of deviations to reach it) between an information set and the equilibrium path. This amounts to saying that players facing more deviations can be expected to have more distant beliefs. In some way, if we note Player 2's payoff changes by dx, and Player 3's payoff changes by dy, it could make sense, in the game in Figure 6, to weight $|\frac{dx}{x}|$ more strongly than $|\frac{dy}{y}|$ in the distance function to minimize (for example $2(\frac{dx}{x})^2$ and $(\frac{dy}{y})^2$). It would perhaps even make more sense to change the power assigned to $|\frac{dx}{x}|$ and $|\frac{dy}{y}|$ given that the probability of several deviations exponentially decreases with the number of deviations (we could work with $(\frac{dx}{x})^2$ and $(\frac{dy}{y})^4$). In this way, given that $|\frac{dx}{x}|$ and

$|\frac{dy}{y}|$ are lower than one, Player 3 can afford more payoff changes without increasing too much the value of the distance function: this amounts to saying that we do not judge his beliefs very distant from the other players', even if in fact they are very different.

5. Revealed beliefs and strategic beliefs

Taking into account the number of deviations to reach an out-of-equilibrium information set puts into light a new problem when studying the notion of distance in beliefs. What exactly is the link between beliefs and deviations? May there be a difference between *announced* beliefs and the beliefs *revealed* through the players' behavior?

Let us again consider the game in Figure 6. Are Player 3's beliefs really distant from Player 2's revealed beliefs? A_2 is optimal when $\mu(x_2) \geq 0.6$. Yet, when Player 3 is called on to play, Player 2 played B_2 and not A_2. Given that Player 2 plays B_2 when his beliefs check $\mu(x_2) < 0.6$ (because the complementary beliefs lead to the play of A_2), we can say that if Player 3 is called on to play (if Player 2 played B_2), then Player 2's *revealed* beliefs contradict the beliefs announced at the information set h. And the revealed beliefs ($\mu(x_2) < 0.6$), are compatible with Player 3's beliefs, $\mu(y_1) = \mu(x_2) \leq 0.3$. So Player 3's beliefs are in reality compatible with Player 2's revealed beliefs. By the way, this might provide a logical reason for the reversal of beliefs of Bonanno's external observer at h' in the game in Figure 6. Given that Player 2 does not play A_2 (the action compatible with the observer's beliefs $\mu(x_2) \geq 0.6$), Player 2 reveals to the observer that his beliefs are the reversed ones, which induces the observer to change his beliefs.[4]

This way of coping with beliefs may seem attractive but it clearly leads to difficulties. First, it is often incompatible with the ICSE beliefs. In the game in Figure 2 for example, Player 3, with the ICSE concept, necessarily has the same beliefs at h' and h'' ($\mu(y_1) = \mu(y_3)$), so his beliefs are not reversed at h'' despite the fact that he knows, at h'', that Player 2 did not behave in conformity with his beliefs at h (given that he did not play A_2). In fact Player 3, at h'', sees Player 2's action B_2 as a trembling hand action that has no informative content and his beliefs only follow from his own perturbations on Player 1's actions A_1 and B_1. Secondly, AGM-consistent beliefs in this game are compatible with the notion of revealed beliefs, given that Player 3's beliefs at h'' take into account that Player 2, by playing B_2, revealed that his beliefs are such

[4] But Bonanno (2013, 2016) does not require this reversal. AGM-consistent assessments also allow beliefs such as $\mu(x_2) = \mu(y_1)$ and $\mu(x_3) = \mu(y_2)$.

that $\mu(x_2) \leq 0.7$. But AGM-consistent beliefs could also assign probability 0.7 to y_3 in that the external observer is in no way compelled to change his view on Player 1's played actions after an unexpected action from Player 2.

Thirdly, taking into account revealed beliefs puts into question the measure of the distance we proposed previously. If a player tries to be close to a previous player's beliefs, then the notion of revealed beliefs requires that his beliefs must be different depending on whether the previous player played the planned action or not. In Figure 2 for example, a SE requires $\mu(x_2) = \mu(y_1) = \mu(y_3)$, so if an equilibrium starts with $\mu(y_1) \neq \mu(y_3)$, the distance is necessarily strictly positive despite the fact that revealed beliefs by definition require $\mu(y_1) \neq \mu(y_3)$. This suggests that taking into account the number of deviations required to reach an information set must change our measure of distance.

Finally, we should consider the notion of revealed beliefs with suspicion. Let us consider the game in Figure 7.

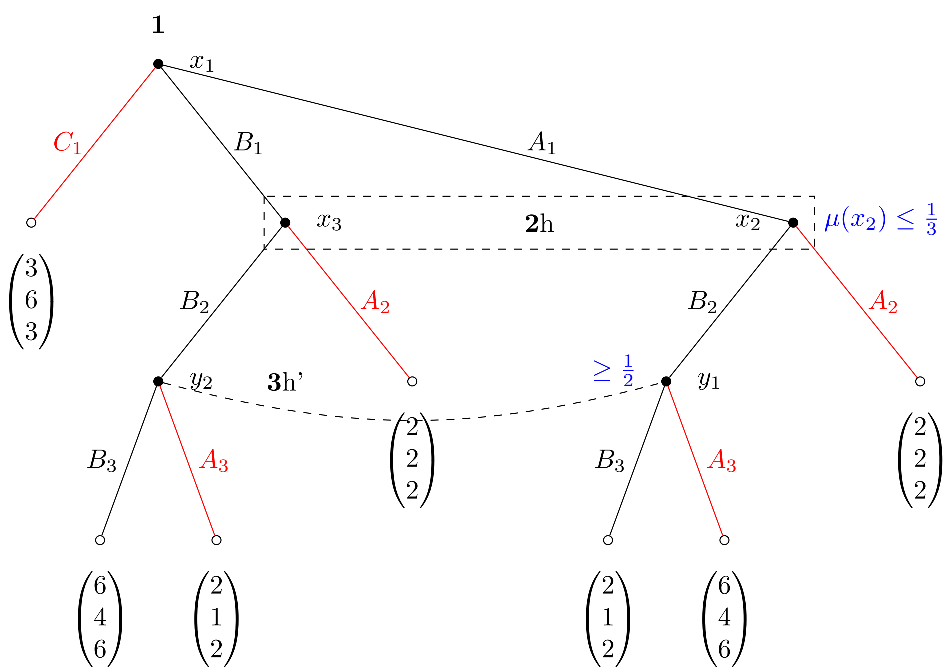

Figure 7: Revealed and announced beliefs.

In every sequential equilibrium, we have $\mu(x_2) = \mu(y_1)$, and it follows that Player 2 always plays B_2 (because either $\mu(x_2) = \mu(y_1) > \frac{1}{2}$, Player 3 plays A_3 and therefore Player 2 plays B_2 or $\mu(x_2) = \mu(y_1) < \frac{1}{2}$, Player 3 plays B_3 and therefore Player 2 plays B_2, or $\mu(x_2) = \mu(y_1) = \frac{1}{2}$, Player 3 plays A_3 with probability q and therefore

Player 2 plays B_2). This implies that Player 1 plays A_1 if $\mu(x_2) = \mu(y_1) > \frac{1}{2}$, B_1 if $\mu(x_2) = \mu(y_1) < \frac{1}{2}$, and A_1 or B_1 depending on the value of q if $\mu(x_2) = \mu(y_1) = \frac{1}{2}$. Hence, both Players 1 and 3 get 6 at any sequential equilibrium and Player 2 gets at best 4.[5]

Now consider the ICSE given in Figure 7, with the SE incompatible beliefs $\mu(x_2) \leq \frac{1}{3}$ and $\mu(y_1) \geq \frac{1}{2}$. This new profile of actions and beliefs checks AGM-consistency (and Bonanno's PBE consistency) and is conform to revealed beliefs. As a matter of fact, given his beliefs and Player 3's action A_3, Player 2 should play A_2, which he does at equilibrium and which induces Player 1 to play C_1. So, if Player 3 is called on to play, this means that Player 2 played B_2 which, according to revealed beliefs, reveals that he does not believe that Player 1 played A_1 with a probability lower than $\frac{1}{3}$. Therefore, Player 3, if he wishes to share similar beliefs to Player 2, can believe that Player 1 played A_1 with a probability larger than $\frac{1}{3}$ and possibly larger than $\frac{1}{2}$. These beliefs induce him to play A_3. So revealed beliefs sustain a profile of strategies where Players 1 and 3 only get 3 and Player 2 gets 6, by leading Player 1 to play C_1.

But let us look more closely at Player 2's revealed beliefs. By playing B_2 instead of the planned action A_2, Player 2 sends the following message to Player 3: "Normally I play A_2 because it is my best response to your action A_3, since I believe that Player 1 much more deviates to B_1 than to A_1; so, if you see me playing B_2, this means that I changed my opinion about Player 1's deviation, so that you are right in believing that he more often deviated to A_1, and so you are right in playing A_3." The problem is that this message is at best "cheap talk" because Player 2 has no information to reveal to Player 3. Player 2 has no idea about the potential deviations of Player 1, given that Player 1 does not deviate when he expects Player 2 to play A_2. One more time, by contrast to Aumann (1976), Geanakoplos and Polemarchakis (1982), and Hart and Tauman (2004), the out-of-equilibrium actions played by the players do not provide information on a true state to discover but on an action that will never be observed at equilibrium. Player 1 does not deviate and so there is nothing to learn about his deviation. If Player 3 "naively" believes that Player 2 can reveal something about Player 1's deviations with his behavior, then he gives Player 2 the power of manipulating Player 3's beliefs to his advantage. As a matter of fact, the beliefs in this ICSE clearly are advantaging Player 2 because they induce Player 1 to play C_1, so they lead to the payoff profile (3,6,3) which is exclusively in advantage of Player 2.

[5] There also exists a SE where Player 1 plays A_1 and B_1 with probability $\frac{1}{2}$ and Player 3 plays A_3 and B_3 with probability $\frac{1}{2}$, so they both get 4 given that Player 2 plays B_2.

6. Discussion

Our analysis suggests that out-of-equilibrium beliefs might be here to justify the players' behaviour at out-of-equilibrium information sets. Choosing them in a given way may help to get a higher payoff, which means that they belong to the strategy set. Let us recall that in a SE, everybody builds the beliefs on a same profile of perturbations, so a player does not really choose his beliefs given that he applies Bayes' Rule to the same perturbations. So in a SE, out-of-equilibrium beliefs do not belong to the player's strategy set. By contrast, in an ICSE, each player chooses his own profile of perturbations, which means that he chooses his own beliefs at out-of-equilibrium sets. This degree of liberty can be exploited to build beliefs that lead to interesting payoffs. In the game in Figure 7, it is good for Player 2 to build beliefs (about Player 1's deviation) that are strikingly different from those of Player 3 in order to justify the action A_2 that prevents Player 1 from deviating from C_1, the most interesting action for Player 2. By contrast, for Player 3, it is better to have similar beliefs than Player 2, in order to lead Player 2 to play B_2, which ultimately leads to the payoffs (6,4,6). So we are tempted to say that, given that each player chooses his set of perturbations, out-of-equilibrium beliefs belong to the strategy set. This fact induces two consequences: we have to reconsider the notion of sequential rationality and we have to reconsider the notion of distance between beliefs.

We start with the notion of distance by coming back to the game studied in Figure 1. We already made the observation, in Section 3, when opting for a proportional approach (Program 2), that rather than asking for $|\frac{dx}{x}| \leq 0.1$ or another threshold, we should ask for a threshold whose value rises with the benefit linked to the ICSE payoffs. As a matter of fact, if everybody benefits from the ICSE payoffs then nobody cares about the distance between the beliefs necessary to sustain the equilibrium. By contrast, in the game in Figure 7, Player 3 may require that Player 2 has beliefs that are close to his own beliefs (especially if Player 2 is a newcomer in the community) because similar beliefs among Player 2 and Player 3 are necessary for Player 3 to get the nice payoff 6. Conversely, if Player 3 is the newcomer in the community, then Player 2 might not pressure Player 3 to adopt beliefs close to his own. The (social) pressure to modify beliefs so as to be close to another player's beliefs must therefore be contrasted with the benefits (in terms of actions played) of holding different beliefs.

Concerning sequential rationality, given that each player chooses his perturbation scheme, these perturbation schemes belong to his strategy set, that is, each player will build (in a consistent way) beliefs at his out-of-equilibrium information sets to get a better payoff. This changes nothing with respect to the definition of individual consistency, but this should lead us to reconsider Kreps and Wilson (1982)'s notion of sequential rationality. In some way, we should add that, for each player, given the strategies played by the other players, there does not exist a perturbation scheme that is sequentially rational (as defined by Kreps and Wilson 1982), individually-consistent and that leads to a larger payoff for the player.

However, sequential rationality rests on unilateral deviations and this additional condition might thus not always help. For example, the ICSE in Figure 7 would resist such an additional condition despite the fact that Player 3 would like to adapt his beliefs to Player 2's beliefs to compel him to play B_2. The problem is that, as long as Player 2 plays A_2 and Player 1 plays C_1, Player 3's beliefs and actions have no impact on his equilibrium payoff. The idea is that each player selects the perturbation scheme associated to the ICSE that leads him to his largest equilibrium payoff. If so, in the game in Figure 7, Player 3 should opt for a perturbation scheme (on Player 1's actions) similar to Player 2's, to push Player 1 and Player 2 to play A_1 and B_2 for example (he can choose the SE (A_1, B_2, A_3) with the beliefs $\mu(x_2) = \mu(y_1) = 1$). Yet, Player 2 would of course choose another perturbation scheme, namely the one leading to the ICSE in Figure 7. So this additional condition, in the game in Figure 7, leads to the non-existence of a system of perturbation schemes both selected by Player 2 and Player 3. In the game in Figure 1, Player 2 and Player 3 can select the same perturbation scheme, namely the one of an ICSE leading to the payoffs (5.99, 10,10). Player 1 can not counter Player 2 and Player 3's selection, in that if they play B_2 and B_3 he is constrained to play C_1. Yet, even in this game, Player 1 may opt for another equilibrium in which he plays A_1, therefore constraining Player 2 and Player 3 to play A_2 and A_3. Therefore, there is no obvious way, even if we switch to coalitions of players, to clearly formalize and express the wish to select payoff-optimizing perturbation schemes. The only trivial configuration is a game such that one ICSE ensures the best payoff to all the players, so that the grand coalition of all the players will be incentivized to select it.

7. Conclusion

The paper started with an obvious observation: there is no reason that leads every player to build the perturbed strategies similarly. Each player has to respect the probabilities assigned to actions that are in the support of the equilibrium, but, given that there does not exist an external observer who can decide for the profiles of ϵ-perturbations, each player is free to build the perturbations assigned to the actions out of the support of the equilibrium. It follows that in an ICSE, players can have different beliefs at out-of-equilibrium information sets.

This led us first to evaluate the distance between different beliefs, because players in the same community are often expected to share similar beliefs. We did this in Section 3 in two ways: (i) an ordering of perturbations and (ii) the minimization of changes in payoffs necessary to make the ICSE beliefs SE-compatible.

We then focused on the function held by beliefs at out-of-equilibrium sets. Since players can build their beliefs at out-of-equilibrium sets, they might build them strategically in order to improve their payoffs. This observation led us to reconsider the traditional concept of sequential rationality, by further requiring that there does not exist a perturbation profile that is individually-consistent and that provides greater payoffs to the player, even though such an additional constraint is not always easy to cope with.

8. Appendix 1

We show that there are no sequential equilibria supporting the action C_1 for player 1 (and therefore the socially optimal situation) in the game shown in Figure 1.

Case 1: Player 2 plays A_2. Therefore player 1 plays A_1.

Case 2: Player 2 plays B_2.

(i): If player 3 plays A_3, player 1 plays B_1.

(ii): If player 3 plays B_3, then player 1 might want to play C_1, but we have already shown that the beliefs that would support this equilibrium are not mutually consistent.

(iii): If player 3 plays A_3 and B_3, then $\gamma = \frac{1}{4}$, and so necessarily $\alpha = \frac{1}{4}$. But then, player 2 would prefer to play A_2. To show this, first let r be the probability that player 3 plays A_3. By playing A_2, player 2's expected payoff is $\frac{1}{4}(2r + 4(1-r) + \frac{3}{4}(4r + 5(1-r)) = \frac{1}{4}(14r + 19(1-r))$. By playing B_2, player 2's expected payoff is $\frac{1}{4}r + \frac{3}{4}(3r + 6(1-r)) = \frac{1}{4}(10r + 18(1-r))$, which is strictly inferior to the expected gain of playing A_2.

Case 3: Player 2 plays A_2 and B_2.

(i): Player 3 plays A_3. In this case, it would not be profitable for player 2 to randomize, given that playing only A_2 would allow her to always gain strictly more.

(ii): Player 3 plays B_3. Given player 2's indifference between A_2 and B_2, it is necessary that $\alpha = \frac{1}{5}$. To show that with these beliefs, player 3 would want to deviate, first note q the probability that player 2 would play A_2. Then by playing A_3, player 3's expected gain would be $\frac{1}{5}(1+2q) + \frac{4}{5}(1+3q) = \frac{1}{5}(5+14q)$. By playing B_3, player 3's expected gain would be $\frac{1}{5}(4-2q) + \frac{4}{5}(3q) = \frac{1}{5}(4+10q)$, which is strictly inferior to the expected gain player 3 would receive by playing A_3.

(iii): Player 3 plays A_3 and B_3. Let r be the probability that player 3 plays A_3, and q the probability that player 2 plays A_2. Let ϵ_0 and ϵ_1 be the perturbations associated to A_1 and B_1 respectively. The expected gain of playing A_2 for player 2 is $\epsilon_0(2r + 4(1-r)) + \epsilon_1(4r + 5(1-r)) = \epsilon_0(4-2r) + \epsilon_1(5-r)$. The expected gain of playing B_2 for player 2 is $\epsilon_0 r + \epsilon_1(3r + 6(1-r)) = \epsilon_0 r + \epsilon_1(6-3r)$. Equalizing these expected gains yields $\epsilon_0(4-2r)+\epsilon_1(5-r) = \epsilon_0 r + \epsilon_1(6-3r)$, or $\epsilon_0(4-3r) = \epsilon_1(1-2r)$.

The expected gain of playing A_3 for player 3 is $\epsilon_0(3q + (1-q)) + \epsilon_1(4q + (1-q)) = \epsilon_0(1+2q) + \epsilon_1(1+3q)$. The expected gain of playing B_3 for player 3 is $\epsilon_0(2q + 4(1-q)) + \epsilon_1 3q = \epsilon_0(4-2q) + \epsilon_1 3q$. Equalizing these expected gains yields $\epsilon_0(1+2q) + \epsilon_1(1+3q) = \epsilon_0(4-2q) + \epsilon_1 3q$, or $\epsilon_0(3-4q) = \epsilon_1$. It follows that $\epsilon_1 = \epsilon_0(3-4q) = \epsilon_0 \frac{4-3r}{1-2r}$. Therefore, $3 - 4q = \frac{4-3r}{1-2r}$, so $4q = 3 - \frac{4-3r}{1-2r} = \frac{3-6r-4+3r}{1-r} = \frac{-1-3r}{1-r}$, which is strictly inferior to 0; an impossible event.

9. Appendix 2

Proof of Proposition 3.1. (i) This follows from the fact that in an ICSE, all strategies are sequentially rational and beliefs are obtained via Bayes' Rule applied to strategies close to the true ones (even if the perturbations are not the same among players). In an ICSE, players agree on the planned actions, even those at unreached subgames, so the ICSE induces a Nash equilibrium in each subgame.

(ii) This follows directly from the definition of both concepts.

(iii) This follows from the fact that the perturbations required by Player 2 (to build his beliefs) are about actions played by Player 1, and the perturbations required by Player 1 are about actions played by Player 2. Both players do not work with different perturbations about actions played by another (third) player. So we can work with one set of perturbations for the game, which is the same for both players (by taking Player 1's perturbations (about Player 2's actions) and Player 2's perturbations (about Player 1's actions)).

(iv) To show why an ICSE is not necessarily a PBE, we choose a game closer to the games studied by Fudenberg and Tirole (1991), by changing our main example in the following way. θ_1 and $\theta_1^{'}$ are Player 1's two possible types, unknown to Player 2 and to Player 3 (prior probabilities ρ and $1-\rho$). So we get the game in Figure 8.

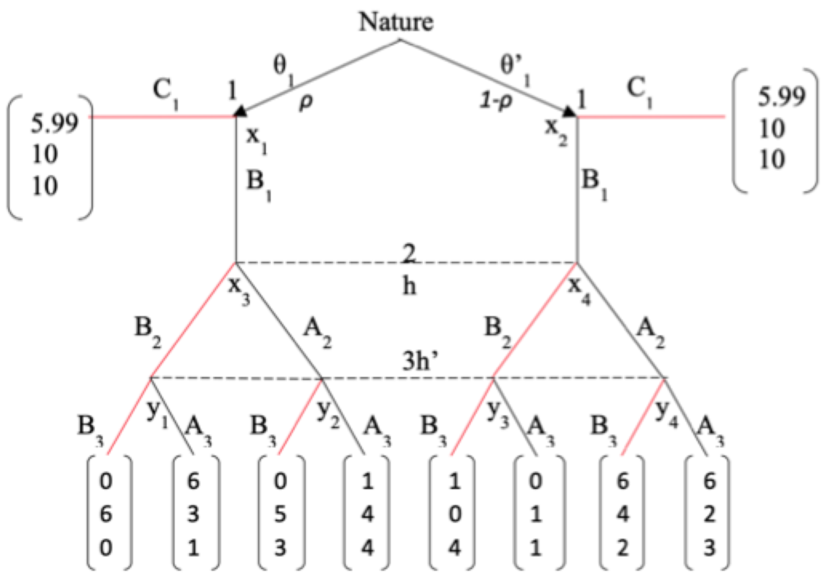

Figure 8: Distinction between ICSE and PBE (1/3).

According to Fudenberg and Tirole (1991)'s PBE equilibrium concept, Player 2 and Player 3 share the same beliefs everywhere (Condition B(iv) p.332), and these beliefs are build using the history of play whenever possible. So, if $\mu(x_4) = \mu_2(\theta_1^{'}/h) = \mu(\theta_1^{'}/h) = \frac{1}{5}$, we get:

$$\begin{aligned}\mu(y_3) &= \frac{\mu(\theta_1^{'}/h)\pi_2(B_2)}{\mu(\theta_1/h)\pi_2(B_2) + \mu(\theta_1/h)\pi_2(A_2) + \mu(\theta_1^{'}/h)\pi_2(B_2) + \mu(\theta_1^{'}/h)\pi_2(A_2)} \\ &= \mu(\theta_1^{'}/h) \\ &= \frac{1}{5},\end{aligned}$$

(due to the Condition B(ii) p.332). Therefore, we can not get $\mu(y_3) = \frac{1}{4}$. Player

3's beliefs are built like Player 2's. Given that B_2 is an expected action, and given that the beliefs at h are not 0, the beliefs at y_3 are necessarily the same than the ones at x_4.

The ICSE concept takes into account that B_2 is an expected action, but it allows Player 3 not to share Player 2's beliefs at h. Therefore, we keep Condition B(ii) but not Condition B(iv), in that Player 2 (respectively Player 3) assigns probability $\frac{1}{5}$ (respectively $\frac{1}{4}$) to θ' if h is reached.

Yet, all PBE are not necessarily ICSE. For example, consider Fudenberg and Tirole (1991)'s example reproduced in Figures 9 and 10 (Figure 8.9 p.346 in their book). The beliefs are in red. The beliefs assigned to the states θ have been obtained by Bayesian inference from previous play. Player 1 is the player who plays the actions a_1^*, $a_1^{'}$ and $a_1^{''}$, while e_k, $e_k^{'}$ are perturbations going to 0.

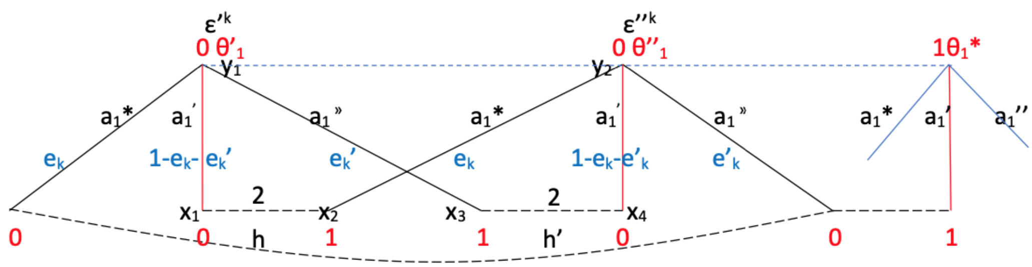

Figure 9: Distinction between ICSE and PBE (2/3).

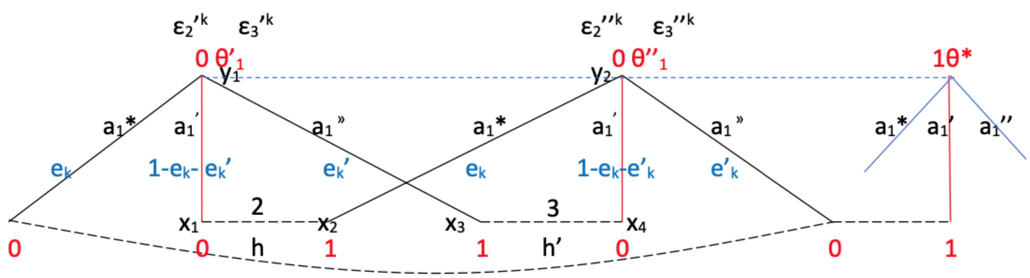

Figure 10: Distinction between ICSE and PBE (3/3).

Fudenberg and Tirole (1991) say that the beliefs (in red) at h and h' belong to a PBE because the PBE places no restrictions on the beliefs at h and h', since both beliefs that lead to h and h' are 0. The only condition is that these beliefs have to be common to all players.

Yet the beliefs at h and h' can not belong to an ICSE when the same player, Player 2 in Figure 9, plays at h and h'. As a matter of fact, by calling ϵ'^k and ϵ''^k the probabilities to reach y_1 and y_2, respectively, which go to 0 given the beliefs, we get $\mu_2(x_3) = \lim\limits_{\epsilon \to 0}(\frac{\epsilon'^k e'_k}{\epsilon'^k e'_k + \epsilon''^k(1-e_k-e_k')})$, which can go to 1 only if $\frac{\epsilon''^k(1-e_k-e_k')}{\epsilon'^k e'_k}$ goes to 0, which requires that $\frac{\epsilon''^k}{\epsilon'^k} \to 0$. But then $\mu_2(x_2) = \lim\limits_{\epsilon \to 0}(\frac{\epsilon''^k e_k}{\epsilon''^k e_k + \epsilon'^k(1-e_k-e_k')}) \to 0$. So the PBE is not a ICSE.

By contrast, when there are two different players at h and h'', like in Figure 10 (Player 2 and Player 3 respectively), the PBE is an ICSE, since we can take different perturbations leading to y_1 and y_2 for Player 2 ($\epsilon_2'^k$ and $\epsilon_2''^k$) and Player 3 ($\epsilon_3'^k$ and $\epsilon_3''^k$), respectively. Therefore, we get: $\mu_3(x_3) = \lim\limits_{\epsilon \to 0}(\frac{\epsilon_3'^k e'_k}{\epsilon_3'^k e'_k + \epsilon_3''^k(1-e_k-e_k')})$, which can go to 1 if $\frac{\epsilon_3''^k}{\epsilon_3'^k e'_k} \to 0$. And we get $\mu_2(x_2) = \lim\limits_{\epsilon \to 0}(\frac{\epsilon_2''^k e_k}{\epsilon_2''^k e_k + \epsilon_2'^k(1-e_k-e_k')})$, which can also go to 1 if $\frac{\epsilon_2'^k}{\epsilon_2''^k e_k} \to 0$. Given that the set of perturbations is different for the two players, both conditions can be fulfilled and the PBE becomes an ICSE.

(v) To formally describe the concept of SCE, we need to introduce some more notation. Let a mixed strategy profile be represented by $\sigma = (\sigma_i)_{i \in N}$, while a behavioral strategy profile will be denoted, as before, by $\pi = (\pi_i)_{i \in N}$. We refer to information sets that are reached with positive probability under the mixed strategy profile σ as $\bar{H}(\sigma)$, and write h_i for a specific information set controlled by player i. We write the behavioral representation of a mixed strategy σ_i as $\hat{\pi}_i(\cdot|\sigma_i)$, such that $\hat{\pi}_i(h_i|\sigma_i)$ represents the probability distribution over actions induced by the mixed strategy σ_i at information set h_i. Furthermore, let μ_i represent player i's beliefs about their opponents' play, such that μ_i is a probability distribution over Π_{-i}, the set of other players' behavioral strategies, with typical element π_{-i}.

A (mixed) strategy profile σ is a *Self-Confirming Equilibrium* (Fudenberg and Levine, 1993) if, $\forall i \in N$ and $\forall s_i \in support(\sigma_i)$, there exists beliefs μ_i such that:

(i) s_i maximizes $u_i(\cdot, \mu_i)$, and

(ii) $\mu_i[\{\pi_{-i} | \pi_j(h_j) = \hat{\pi}_j(h_j|\sigma_j)\}] = 1, \forall j \neq i, \forall h_j \in \bar{H}(s_i, \sigma_{-i})$.

In words, every players' subjective probability distribution needs to put probability 1 to strategy profiles that are compatible with observed play (reached with positive probability). That is, players' expectations need to be right on the equilibrium path, but need not be right at information sets that are never reached.

Importantly, since each player has to best respond only to the *observed* actions of other players, the SCE only requires that players play a best response to their *beliefs* about other player's actions out of the equilibrium strategy path. This implies that a SCE is not necessarily an ISCE, since the ICSE requires that players play best responses to other player's *actions*, even at information sets that are out of the equilibrium strategy path. On the other hand, an ICSE is always a SCE, since it is always possible to create a SCE in which players play best responses, and have accurate beliefs, even at out-of-equilibrium information sets.

■

Proof of Proposition 3.2. (i) AGM-consistency works with plausibility orders on stories. It is a qualitative notion that focuses on plausibility-preserving actions, which are actions played with a positive probability in the game. Given that the ICSE also respects the actions played with a positive probability in the game, an ICSE is usually AGM-consistent, except if it leads to 0-1 assessments.

(ii) Bonanno's PBE concept transforms the qualitative notion into a quantitative one, by requiring that the beliefs respect Bayes' Rule when it applies, i.e., in presence of plausibility-preserving actions. Given that Bonanno's concept only introduces one probability distribution at each out-of-equilibrium information set, it immediately follows that it cannot intersect with the ICSE concept, which works with the Bayesian rule applied to several perturbation distributions, one for each player.

■

10. Appendix 3

Program 1:

$$\min_{x_1,x_2,x_3,x_4,y_1,y_2,y_3,y_4,\mu} \quad {x_1}^2+{x_2}^2+{x_3}^2+{x_4}^2+{y_1}^2+{y_2}^2+{y_3}^2+{y_4}^2$$

$$\begin{aligned} \text{s.t.} \quad & (0+x_2)\mu+(6+x_4)(1-\mu) \geq (4+x_1)\mu+(5+x_3)(1-\mu) \quad \mathbf{(1)}\ \lambda_1 \\ & (4+y_2)\mu+(0+y_4)(1-\mu) \geq (1+y_1)\mu+(1+y_3)(1-\mu) \quad \mathbf{(2)}\ \lambda_2 \\ & \mu \geq 0 \quad \mathbf{(3)}\ \lambda_3 \\ & 1-\mu \geq 0 \quad \mathbf{(4)}\ \lambda_4 \end{aligned}$$

Equations **(1)** and **(2)** can be rewritten:

$$1 - x_3 + x_4 - \mu(5 + x_1 - x_2 - x_3 + x_4) \geq 0 \quad \textbf{(1)}$$
$$-1 - y_3 + y_4 - \mu(-4 + y_1 - y_2 - y_3 + y_4) \geq 0 \quad \textbf{(2)}$$

The KT function becomes:

$$x_1{}^2 + x_2{}^2 + x_3{}^2 + x_4{}^2 + y_1{}^2 + y_2{}^2 + y_3{}^2 + y_4{}^2 - \lambda_1(1 - x_3 + x_4 - \mu(5 + x_1 - x_2 - x_3 + x_4))$$
$$- \lambda_2(-1 - y_3 + y_4 - \mu(-4 + y_1 - y_2 - y_3 + y_4)) - \lambda_3\mu - \lambda_4(1 - \mu).$$

The KT equations are:

$$2x_1 + \lambda_1\mu = 0 \quad \textbf{(a)}$$
$$2x_2 - \lambda_1\mu = 0 \quad \textbf{(b)}$$
$$2x_3 + \lambda_1(1 - \mu) = 0 \quad \textbf{(c)}$$
$$2x_4 - \lambda_1(1 - \mu) = 0 \quad \textbf{(d)}$$
$$2y_1 + \lambda_2\mu = 0 \quad \textbf{(e)}$$
$$2y_2 - \lambda_2\mu = 0 \quad \textbf{(f)}$$
$$2y_3 + \lambda_2(1 - \mu) = 0 \quad \textbf{(g)}$$
$$2y_4 - \lambda_2(1 - \mu) = 0 \quad \textbf{(h)}$$
$$\lambda_1(5 + x_1 - x_2 - x_3 + x_4) + \lambda_2(-4 + y_1 - y_2 - y_3 + y_4) - \lambda_3 + \lambda_4 = 0 \quad \textbf{(i)}$$

It follows that $x_2 = -x_1 \geq 0$, $x_4 = -x_3 \geq 0$, $y_2 = -y_1 \geq 0$, and $y_4 = -y_3 \geq 0$, due to the positivity of the KT multipliers.

Both Conditions **(1)** and **(2)** are necessarily checked with equality, because $\lambda_1 = 0$ leads to $\lambda_2(-4 - 2y_2 + 2y_4) = 0$, hence $y_4 \geq 2$, which can clearly not lead to a global minimum given our numerical introduction, and $\lambda_2 = 0$ leads to $\lambda_1(5 - 2x_2 + 2x_4) = 0$, hence $x_4 \geq 2.5$, which can not lead to a global minimum for the same reason.

We seek a solution that checks $0 < \mu < 1$, and so $\lambda_3 = \lambda_4 = 0$. It follows that $\lambda_1 = \frac{2x_2}{\mu} = \frac{2x_4}{1-\mu}$, hence $\mu = \frac{x_2}{x_2+x_4}$, and $\lambda_2 = \frac{2y_2}{\mu} = \frac{2y_4}{1-\mu}$, hence $\mu = \frac{y_2}{y_2+y_4}$. As a result:

$$x_2 y_4 = y_2 x_4 \quad \textbf{(j)}$$
$$2(x_2^2 + x_4^2) = 4x_2 - x_4 \quad \textbf{(k)}$$
$$2(y_2^2 + y_4^2) = -3y_2 + y_4 \quad \textbf{(l)}$$
$$x_2(1 + 2x_4) + y_2(-1 + 2y_4) = 0 \quad \textbf{(m)}$$

The only solution is:

$$
\begin{aligned}
x_2 = -x_1 &= -\frac{53\sqrt{29}-290}{290} \simeq 0.0158\\
x_4 = -x_3 &= \frac{33\sqrt{29}-145}{580} \simeq 0.0564\\
y_2 = -y_1 &= \frac{83\sqrt{29}-435}{580} \simeq 0.0206\\
y_4 = -y_3 &= -\frac{19\sqrt{29}-145}{580} \simeq 0.0736\\
\mu &= \frac{-106\sqrt{29}+580}{435-73\sqrt{29}} \simeq 0.219\\
\lambda_1 &= \frac{435-73\sqrt{29}}{290} \simeq 0.1444\\
\lambda_2 &= \frac{64\sqrt{29}-290}{290} \simeq 0.1885.
\end{aligned}
$$

Program 2:

$$
\begin{aligned}
\min_{x_1,x_2,x_3,y_1,y_2,y_3,\mu} \quad & (\frac{x_1}{4})^2 + (\frac{x_2}{5})^2 + (\frac{x_3}{6})^2 + {y_1}^2 + (\frac{y_2}{4})^2 + {y_3}^2\\
\text{s.t.} \quad & (6+x_3)(1-\mu) \geq (4+x_1)\mu + (5+x_2)(1-\mu) \quad \mathbf{(1)}\ \lambda_1\\
& (4+y_2)\mu \geq (1+y_1)\mu + (1+y_3)(1-\mu) \quad \mathbf{(2)}\ \ \lambda_2\\
& \mu \geq 0 \quad \mathbf{(3)}\ \lambda_3\\
& 1-\mu \geq 0 \quad \mathbf{(4)}\ \lambda_4
\end{aligned}
$$

Equations **(1)** and **(2)** can be rewritten:

$$
\begin{aligned}
1 + x_3 - x_2 + \mu(-5 - x_1 + x_2 - x_3) \geq 0 \quad \mathbf{(1)}\\
-1 - y_3 + \mu(4 - y_1 + y_2 + y_3) \geq 0 \quad \mathbf{(2)}
\end{aligned}
$$

The KT function becomes:

$$
\begin{aligned}
&(\frac{x_1}{4})^2 + (\frac{x_2}{5})^2 + (\frac{x_3}{6})^2 + {y_1}^2 + (\frac{y_2}{4})^2 + {y_3}^2 - \lambda_1(1 + x_3 - x_2 + \mu(-5 - x_1 + x_2 - x_3))\\
&- \lambda_2(-1 - y_3 + \mu(4 - y_1 + y_2 + y_3)) - \lambda_3\mu - \lambda_4(1-\mu).
\end{aligned}
$$

The KT equations are:

$$\frac{2x_1}{16} + \lambda_1\mu = 0 \quad \textbf{(a)}$$
$$\frac{2x_2}{25} + \lambda_1(1-\mu) = 0 \quad \textbf{(b)}$$
$$\frac{2x_3}{36} - \lambda_1(1-\mu) = 0 \quad \textbf{(c)}$$
$$2y_1 + \lambda_2\mu = 0 \quad \textbf{(e)}$$
$$\frac{2y_2}{16} - \lambda_2\mu = 0 \quad \textbf{(f)}$$
$$2y_3 + \lambda_2(1-\mu) = 0 \quad \textbf{(g)}$$
$$\lambda_1(5 + x_1 - x_2 + x_3) + \lambda_2(-4 + y_1 - y_2 - y_3) - \lambda_3 + \lambda_4 = 0 \quad \textbf{(i)}$$

We look for a solution such that Conditions **(1)** and **(2)** are checked with equality, and such that $0 < \mu < 1$ (so that $\lambda_3 = \lambda_4 = 0$. The only solution gives: $x_1 = -0.0258$, $x_2 = -0.1232$, $x_3 = 0.1774$, $y_1 = -0.0021$, $y_2 = 0.0338$, $y_3 = -0.00645$, $\lambda_1 = 0.01308$, $\lambda_2 = 0.01712$ and $\mu = 0.24657$.

References

Aumann, R. J. (1976): "Agreeing to disagree," *The annals of statistics*, 4, 1236–1239.

Barber, M. and J. C. Pope (2019): "Does party trump ideology? Disentangling party and ideology in America," *American Political Science Review*, 113, 38–54.

Bonanno, G. (2013): "AGM-consistency and perfect Bayesian equilibrium. Part I: definition and properties," *International Journal of Game Theory*, 42, 567–592.

——— (2016): "AGM-consistency and perfect Bayesian equilibrium. Part II: from PBE to sequential equilibrium," *International Journal of Game Theory*, 45, 1071–1094.

Fudenberg, D. and D. K. Levine (1993): "Self-confirming equilibrium," *Econometrica: Journal of the Econometric Society*, 523–545.

Fudenberg, D. and J. Tirole (1991): *Game theory*, MIT press.

Geanakoplos, J. D. and H. M. Polemarchakis (1982): "We can't disagree forever," *Journal of Economic theory*, 28, 192–200.

Gould, E. D. and E. F. Klor (2019): "Party hacks and true believers: The effect of party affiliation on political preferences," *Journal of Comparative Economics*, 47, 504–524.

Greenberg, J., S. Gupta, and X. Luo (2009): "Mutually acceptable courses of action," *Economic Theory*, 40, 91–112.

Hart, S. and Y. Tauman (2004): "Market crashes without external shocks," *The Journal of Business*, 77, 1–8.

Kreps, D. M. and R. Wilson (1982): "Sequential equilibria," *Econometrica: Journal of the Econometric Society*, 863–894.

Slothuus, R. and M. Bisgaard (2021): "How political parties shape public opinion in the real world," *American Journal of Political Science*, 65, 896–911.

Umbhauer, G. and A. Wolff (2019): "Individually-Consistent Sequential Equilibrium," *BETA Working Paper*.