

Game Theory with Translucent Players

Joseph Y. Halpern and Rafael Pass*
Cornell University
Department of Computer Science
Ithaca, NY, 14853, U.S.A.
e-mail: {halpern,rafael}@cs.cornell.edu

July 23, 2018

Abstract

A traditional assumption in game theory is that players are opaque to one another—if a player changes strategies, then this change in strategies does not affect the choice of other players’ strategies. In many situations this is an unrealistic assumption. We develop a framework for reasoning about games where the players may be *translucent* to one another; in particular, a player may believe that if she were to change strategies, then the other player would also change strategies. Translucent players may achieve significantly more efficient outcomes than opaque ones.

Our main result is a characterization of strategies consistent with appropriate analogues of common belief of rationality. *Common Counterfactual Belief of Rationality (CCBR)* holds if (1) everyone is rational, (2) everyone counterfactually believes that everyone else is rational (i.e., all players i believe that everyone else would still be rational even if i were to switch strategies), (3) everyone counterfactually believes that everyone else is rational, and counterfactually believes that everyone else is rational, and so on. CCBR characterizes the set of strategies surviving iterated removal of *minimax dominated* strategies, where a strategy σ for player i is minimax dominated by σ' if the worst-case payoff for i using σ' is better than the best possible payoff using σ .

*Halpern is supported in part by NSF grants IIS-0812045, IIS-0911036, and CCF-1214844, by AFOSR grant FA9550-08-1-0266, and by ARO grant W911NF-09-1-0281. Pass is supported in part by a Alfred P. Sloan Fellowship, Microsoft New Faculty Fellowship, NSF Award CNS-1217821, NSF CAREER Award CCF-0746990, NSF Award CCF-1214844, AFOSR YIP Award FA9550-10-1-0093, and DARPA and AFRL under contract FA8750-11-2-0211. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government.

1 Introduction

Two large firms 1 and 2 need to decide whether to *cooperate* (C) or *sue* (S) the other firm. Suing the other firm always has a small positive reward, but being sued induces a high penalty p ; more precisely, $u(C, C) = (0, 0)$; $u(C, S) = (-p, r)$; $u(S, C) = (r, -p)$, $u(S, S) = (r - p, r - p)$. In other words, we are considering an instance of the Prisoner's Dilemma.

But there is a catch. Before acting, each firm needs to discuss their decision with its board. Although these discussions are held behind closed doors, there is always the possibility of the decision being “leaked”; as a consequence, the other company may change its course of action. Furthermore, both companies are aware of this fact. In other words, the players are *translucent* to one another.

In such a scenario, it may well be rational for both companies to cooperate. For instance, consider the following situation.

- Firm i believes that its action is leaked to firm $2 - i$ with probability ϵ .
- Firm i believes that if the other firm $2 - i$ finds out that i is defecting, then $2 - i$ will also defect.
- Finally, $p\epsilon > r$ (i.e., the penalty for being sued is significantly higher than the reward of suing the other company).

Neither firm defects, since defection is noticed by the other firm with probability ϵ , which (according to their beliefs) leads to a harsh punishment. Thus, the possibility of the players' actions being leaked to the other player allows the players to significantly improve social welfare in equilibrium. (This suggests that it may be mutually beneficial for two countries to spy on each other!)

Even if the Prisoner's dilemma is not played by corporations but by individuals, each player may believe that if he chooses to defect, his “guilt” over defecting may be visible to the other player. (Indeed, facial and bodily cues such as increased pupil size are often associated with deception; see e.g., [Ekman and Friesen 1969].) Thus, again, the players may choose to cooperate out of fear that if they defect, the other player may detect it and act on it.

Our goal is to capture this type of reasoning formally. We take a Bayesian approach: Each player has a (subjective) probability distribution (describing the player's beliefs) over the states of the world. Traditionally, a player i is said to be rational in a state ω if the strategy σ_i that i plays at ω is a best response to the strategy profile μ_{-i} of the other players induced by i 's beliefs in ω ;¹ that is,

¹Formally, we assume that i has a distribution on states, and at each state, a pure strategy profile is played; the distribution on states clearly induces a distribution on strategy profiles for the players other than i , which we denote μ_{-i} .

$u_i(\sigma_i, \mu_{-i}) \geq u_i(\sigma'_i, \mu_{-i})$ for all alternative strategies σ'_i for i . In our setting, things are more subtle. Player i may believe that if she were to switch strategies from σ_i to σ'_i , then players other than i might also switch strategies. We capture this using *counterfactuals* [Lewis 1973; Stalnaker 1968].² Associated with each state of the world ω , each player i , and $f(\omega, i, \sigma'_i)$ where player i plays σ'_i . Note that if i changes strategies, then this change in strategies may start a chain reaction, leading to further changes. We can think of $f(\omega, i, \sigma'_i)$ as the steady-state outcome of this process: the state that would result if i switched strategies to σ'_i . Let $\mu_{f(\omega, i, \sigma'_i)}$ be the distribution on strategy profiles of $-i$ (the players other than i) induced by i 's beliefs at ω about the steady-state outcome of this process. We say that i is rational at a state ω where i plays σ_i and has beliefs μ_i if $u_i(\sigma_i, \mu_{-i}) \geq u_i(\sigma'_i, \mu_{f(\omega, i, \sigma'_i)})$ for every alternative strategy σ'_i for i . Note that we have required the closest-state function to be deterministic, returning a unique state, rather than a distribution over states. While this may seem incompatible with the motivating scenario, it does not seem so implausible in our context that, by taking a rich enough representation of states, we can assume that a state contains enough information about players to resolve uncertainty about what strategies they would use if one player were to switch strategies.

We are interested in considering analogues to rationalizability in a setting with translucent players, and providing epistemic characterizations of them. To do that, we need some definitions. We say that a player i *counterfactually believes* φ at ω if i believes φ holds even if i were to switch strategies. *Common Counterfactual Belief of Rationality (CCBR)* holds if (1) everyone is rational, (2) everyone counterfactually believes that everyone else is rational (i.e., all players i believe that everyone else would still be still rational even if i were to switch strategies), (3) everyone counterfactually believes that everyone else is rational, and counterfactually believes that everyone else is rational, and so on.

Our main result is a characterization of strategies consistent with CCBR. Roughly

²A different, more direct, approach for capturing our original motivating example would be to consider and analyze an extensive-form variant G' of the original normal-form game G that explicitly models the “leakage” of players’ actions in G , allows the player to react to these leakage signals by choosing a new action in G , which again may be leaked and the players may react to, and so on. Doing this is subtle. We would need to model how players respond to receiving leaked information, and to believing that there was a change in plan even if information wasn’t leaked. To make matters worse, it’s not clear what it would mean that a player is “intending” to perform an action a if players can revise what they do as the result of a leak. Does it mean that a player will do a if no information is leaked to him? What if no information is leaked, but he believes that the other side is planning to change their plans in any case? In addition, modeling the game in this way would require a distribution over leakage signals to be exogenously given (as part of the description of the game G'). Moreover, player strategies would have to be infinite objects, since there is no bound on the sequence of leaks and responses to leaks. In contrast, using counterfactuals, we can directly reason about the original (finite) game G .

speaking, these results can be summarized as follows:

- If the closest-state function respects “unilateral deviations”—when i switches strategies, the strategies and beliefs of players other than i remain the same—then CCBR characterizes the set of rationalizable strategies.
- If the closest-state function can be arbitrary, CCBR characterizes the set of strategies that survive iterated removal of *minimax dominated* strategies: a strategy σ_i is minimax dominated for i if there exists a strategy σ'_i for i such that $\min_{\mu'_{-i}} u_i(\sigma'_i, \mu'_{-i}) > \max_{\mu_{-i}} u_i(\sigma_i, \mu_{-i})$; that is, $u_i(\sigma'_i, \mu'_{-i}) > u_i(\sigma_i, \mu_{-i})$ no matter what the strategy profiles μ_{-i} and μ'_{-i} are.

We also consider analogues of Nash equilibrium in our setting, and show that individually rational strategy profiles that survive iterated removal of minimax dominated strategies characterize such equilibria.

Note that in our approach, each player i has a *belief* about how the other players’ strategies would change if i were to change strategies, but we do not require i to explicitly specify how he would respond to other people changing strategies. The latter approach, of having each player specify how she responds to her opponents’ actions, goes back to von Neumann and Morgenstern [von Neumann and Morgenstern 1947, pp. 105–106]:

Indeed, the rules of the game Γ prescribe that each player must make his choice (his personal move) in ignorance of the outcome of the choice of his adversary. It is nevertheless conceivable that one of the players, say 2, ” finds out”; i.e., has somehow acquired the knowledge as to what his adversary’s strategy is. The basis for this knowledge does not concern us; it may (but need not) be experience from previous plays.

Von Neumann and Morgenstern’s analysis corresponds to a *single* round of removal of minimax dominated strategies. This approach was further explored and formalized by Howard [1971] in the 1960s. In Howard’s approach, players pick a “meta-strategy” that takes as input the strategy of other players. It led to complex formalisms involving infinite hierarchies of meta-strategies: at the lowest level, each player specifies a strategy in the original game; at level k , each player specifies a “response rule” (i.e., a meta-strategy) to other players’ $(k - 1)$ -level response rules. Such hierarchical structures have not proven useful when dealing with applications. Since we do not require players to specify reaction rules, we avoid the complexities of this approach.

Program equilibria [Tennenholz 2004] and *conditional commitments* [Kalai, Kalai, Lehrer, and Samet 2010] provide a different approach to avoiding infinite hierarchies. Roughly speaking,

each player i simply specifies a *program* Π_i ; player i 's action is determined by running i 's program on input the (description of) the programs of the other players; that is, i ' action is given by $\Pi_i(\Pi_{-i})$. Tennenholtz [2004] and Kalai et al. [2010] show that every (correlated) individually rational outcome can be sustained in a program equilibrium. Their model, however, assumes that player's programs (which should be interpreted as their "plan of action") are commonly known to all players. We dispense with this assumption. It is also not clear how to define common belief of rationality in their model; the study of program equilibria and conditional commitments has considered only analogues of Nash equilibrium.

Perhaps most closely related to our model is a paper by Spohn [2003] that studies a generalization of Nash equilibrium called *dependency equilibria*, where players' conjectures are described as "conditional probabilities": for each action a_1 of player 1, player 1 may have a different belief about the action of player 2. Independently of our work, Salcedo [2013], defines a notion of *conjectural rationalizability* that replaces beliefs (over actions) by conjectures described as conditional probabilities, as in [Spohn 2003]. Salcedo also defines a notion of minimax domination (which he calls *absolute domination*), and characterizes conjectural rationalizability in terms of strategies surviving iterated deletion of minimax dominated strategies.

Counterfactuals have been explored in a game-theoretic setting; see, for example, [Aumann 1995; Halpern 1999; Samet 1996; Stalnaker 1996; Zambrano 2004]. However, all these papers considered only structures where, in the closest state where i changes strategies, all other players' strategies remain the same; thus, these approaches are not applicable in our context.

2 Counterfactual Structures

Given a game Γ , let $\Sigma_i(\Gamma)$ denote player i 's pure strategies in Γ (we occasionally omit the parenthetical Γ if it is clear from context or irrelevant).

To reason about the game Γ , we consider a class of Kripke structures corresponding to Γ . For simplicity, we here focus on finite structures. A *finite probability structure* M appropriate for Γ is a tuple $(\Omega, \mathbf{s}, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$, where Ω is a finite set of states; \mathbf{s} associates with each state $\omega \in \Omega$ a pure strategy profile $\mathbf{s}(\omega)$ in the game Γ ; and, for each player i , \mathcal{PR}_i is a *probability assignment* that associates with each state $\omega \in \Omega$ a probability distribution $\mathcal{PR}_i(\omega)$ on Ω , such that

1. $\mathcal{PR}_i(\omega)(\llbracket \mathbf{s}_i(\omega) \rrbracket_M) = 1$, where for each strategy σ_i for player i , $\llbracket \sigma_i \rrbracket_M = \{\omega : \mathbf{s}_i(\omega) = \sigma_i\}$, where $\mathbf{s}_i(\omega)$ denotes player i 's strategy in the strategy profile $\mathbf{s}(\omega)$;

2. $\mathcal{PR}_i(\omega)(\llbracket \mathcal{PR}_i(\omega), i \rrbracket_M) = 1$, where for each probability measure π and player i , $\llbracket \pi, i \rrbracket_M = \{\omega : \mathcal{PR}_i(\omega) = \pi\}$.

These assumptions say that player i assigns probability 1 to his actual strategy and beliefs.

To deal with counterfactuals, we augment probability structures with a “closest-state” function f that associates with each state ω , player i , and strategy σ'_i , a state $f(\omega, i, \sigma'_i)$ where player i plays σ'_i ; if σ'_i is already played in ω , then the closest state to ω where σ'_i is played is ω itself. Formally, a *finite counterfactual structure* M appropriate for Γ is a tuple $(\Omega, \mathbf{s}, f, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$, where $(\Omega, \mathbf{s}, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$ is a probability structure appropriate for Γ and f is a “closest-state” function. We require that if $f(\omega, i, \sigma'_i) = \omega'$, then

1. $\mathbf{s}_i(\omega') = \sigma'_i$;
2. if $\sigma'_i = \mathbf{s}_i(\omega)$, then $\omega' = \omega$.

Given a probability assignment \mathcal{PR}_i for player i , we define i ’s counterfactual belief at state ω (“what i believes would happen if he switched to σ'_i at ω) as

$$\mathcal{PR}_{i, \sigma'_i}^c(\omega)(\omega') = \sum_{\{\omega'' \in \Omega : f(\omega'', i, \sigma'_i) = \omega'\}} \mathcal{PR}_i(\omega)(\omega'').$$

Note that the conditions above imply that each player i knows what strategy he would play if he were to switch; that is, $\mathcal{PR}_{i, \sigma'_i}^c(\omega)(\llbracket \sigma'_i \rrbracket_M) = 1$.

Let $\text{Supp}(\pi)$ denote the support of the probability measure π . Note that $\text{Supp}(\mathcal{PR}_{i, \sigma'_i}^c(\omega)) = \{f(\omega', i, \sigma'_i) : \omega' \in \text{Supp}(\mathcal{PR}_i(\omega))\}$. Moreover, it is almost immediate from the definition that if $\mathcal{PR}_i(\omega) = \mathcal{PR}_i(\omega')$, then $\mathcal{PR}_{i, \sigma'_i}^c(\omega) = \mathcal{PR}_{i, \sigma'_i}^c(\omega')$ for all strategies σ'_i for player i . But it does *not* in general follow that i knows his counterfactual beliefs at ω , that is, it may not be the case that for all strategies σ'_i for player i , $\mathcal{PR}_{i, \sigma'_i}^c(\omega)(\llbracket \mathcal{PR}_{i, \sigma'_i}^c(\omega), i \rrbracket_M) = 1$. Suppose that we think of a state as representing each player’s *ex ante* view of the game. The fact that player $\mathbf{s}_i(\omega) = \sigma_i$ should then be interpreted as “ i intends to play σ_i at state ω .” With this view, suppose that ω is a state where $\mathbf{s}_i(\omega)$ is a conservative strategy, while σ'_i is a rather reckless strategy. It seems reasonable to expect that i ’s subjective beliefs regarding the likelihood of various outcomes may depend in part on whether he is in a conservative or reckless frame of mind. We can think of $\mathcal{PR}_{i, \sigma'_i}^c(\omega)(\omega')$ as the probability that i ascribes, at state ω , to ω' being the outcome of i switching to strategy σ'_i ; thus, $\mathcal{PR}_{i, \sigma'_i}^c(\omega)(\omega')$ represents i ’s evaluation of the likelihood of ω' when he is in a conservative frame of mind. This may not be the evaluation that i uses in states in the support $\mathcal{PR}_{i, \sigma'_i}^c(\omega)$; at all these states, i is in a “reckless” frame of mind. Moreover,

there may not be a unique reckless frame of mind, so that i may not have the same beliefs at all the states in the support of $\mathcal{PR}_{i,\sigma'_i}^c(\omega)$.

M is a *strongly appropriate counterfactual structure* if it is an appropriate counterfactual structure and, at every state ω , every player i knows his counterfactual beliefs. As the example above suggests, strong appropriateness is a nontrivial requirement. As we shall see, however, our characterization results hold in both appropriate and strongly appropriate counterfactual structures.

Note that even in strongly appropriate counterfactual structures, we may not have $\mathcal{PR}_i(f(\omega, i, \sigma'_i)) = \mathcal{PR}_{i,\sigma'_i}^c(\omega)$. We do have $\mathcal{PR}_i(f(\omega, i, \sigma'_i)) = \mathcal{PR}_{i,\sigma'_i}^c(\omega)$ in strongly appropriate counterfactual structures if $f(\omega, i, \sigma'_i)$ is in the support of $\mathcal{PR}_{i,\sigma'_i}^c(\omega)$ (which will certainly be the case if ω is in the support of $\mathcal{PR}_i(\omega)$). To see why we may not want to have $\mathcal{PR}_i(f(\omega, i, \sigma'_i)) = \mathcal{PR}_{i,\sigma'_i}^c(\omega)$ in general, even in strongly appropriate counterfactual structures, consider the example above again. Suppose that, in state ω , although i does not realize it, he has been given a drug that affects how he evaluates the state. He thus ascribes probability 0 to ω . In $f(\omega, i, \sigma'_i)$ he has also been given the drug, and the drug in particular affects how he evaluates outcomes. Thus, i 's beliefs in the state $f(\omega, i, \sigma'_i)$ are quite different from his beliefs in all states in the support of $\mathcal{PR}_{i,\sigma'_i}^c(\omega)$.

2.1 Logics for Counterfactual Games

Let $\mathbb{L}(\Gamma)$ be the language where we start with *true* and the primitive proposition RAT_i and $play_i(\sigma_i)$ for $\sigma_i \in \Sigma_i(\Gamma)$, and close off under the modal operators B_i (player i believes) and B_i^* (player i counterfactually believes) for $i = 1, \dots, n$, CB (common belief), and CB^* (common counterfactual belief), conjunction, and negation. We think of $B_i\varphi$ as saying that “ i believes φ holds with probability 1” and $B_i^*\varphi$ as saying “ i believes that φ holds with probability 1, even if i were to switch strategies”.

Let \mathbb{L}^0 be defined exactly like \mathbb{L} except that we exclude the “counterfactual” modal operators B^* and CB^* . We first define semantics for \mathbb{L}^0 using probability structures (without counterfactuals). We define the notion of a formula φ being true at a state ω in a probability structure M (written $(M, \omega) \models \varphi$) in the standard way, by induction on the structure of φ , as follows:

- $(M, \omega) \models \text{true}$ (so *true* is vacuously true).
- $(M, \omega) \models play_i(\sigma_i)$ iff $\sigma_i = s_i(\omega)$.
- $(M, \omega) \models \neg\varphi$ iff $(M, \omega) \not\models \varphi$.
- $(M, \omega) \models \varphi \wedge \varphi'$ iff $(M, \omega) \models \varphi$ and $(M, \omega) \models \varphi'$.

- $(M, \omega) \models B_i \varphi$ iff $\mathcal{PR}_i(\omega)(\llbracket \varphi \rrbracket_M) = 1$, where $\llbracket \varphi \rrbracket_M = \{\omega : (M, \omega) \models \varphi\}$.
- $(M, \omega) \models RAT_i$ iff $\mathbf{s}_i(\omega)$ is a best response given player i 's beliefs regarding the strategies of other players induced by \mathcal{PR}_i .
- Let $EB\varphi$ ("everyone believes φ ") be an abbreviation of $B_1\varphi \wedge \dots \wedge B_n\varphi$; and define $EB^k\varphi$ for all k inductively, by taking $EB^1\varphi$ to be $EB\varphi$ and $EB^{k+1}\varphi$ to be $EB(EB^k\varphi)$.
- $(M, \omega) \models CB\varphi$ iff $(M, \omega) \models EB^k\varphi$ for all $k \geq 1$.

Semantics for \mathbb{L}^0 in counterfactual structures is defined in an identical way, except that we redefine RAT_i to take into account the fact that player i 's beliefs about the strategies of players $-i$ may change if i changes strategies.

- $(M, \omega) \models RAT_i$ iff for every strategy σ'_i for player i ,

$$\sum_{\omega' \in \Omega} \mathcal{PR}_i(\omega)(\omega') u_i(\mathbf{s}_i(\omega), \mathbf{s}_{-i}(\omega')) \geq \sum_{\omega' \in \Omega} \mathcal{PR}_{i, \sigma'_i}^c(\omega)(\omega') u_i(\sigma'_i, \mathbf{s}_{-i}(\omega')).$$

The condition above is equivalent to requiring that

$$\sum_{\omega' \in \Omega} \mathcal{PR}_i(\omega)(\omega') u_i(\mathbf{s}_i(\omega), \mathbf{s}_{-i}(\omega')) \geq \sum_{\omega' \in \Omega} \mathcal{PR}_i(\omega)(\omega') u_i(\sigma'_i, \mathbf{s}_{-i}(f(\omega', i, \sigma'_i))).$$

Note that, in general, this condition is different from requiring that $\mathbf{s}_i(\omega)$ is a best response given player i 's beliefs regarding the strategies of other players induced by \mathcal{PR}_i .

To give the semantics for \mathbb{L} in counterfactual structures, we now also need to define the semantics of B_i^* and CB^* :

- $(M, \omega) \models B_i^* \varphi$ iff for all strategies $\sigma'_i \in \Sigma_i(\Gamma)$, $\mathcal{PR}_{i, \sigma'_i}^c(\omega)(\llbracket \varphi \rrbracket_M) = 1$.
- $(M, \omega) \models CB^* \varphi$ iff $(M, \omega) \models (EB^*)^k \varphi$ for all $k \geq 1$.

It is easy to see that, like B_i , B_i^* depends only on i 's beliefs; as we observed above, if $\mathcal{PR}_i(\omega) = \mathcal{PR}_i(\omega')$, then $\mathcal{PR}_{i, \sigma'_i}^c(\omega) = \mathcal{PR}_{i, \sigma'_i}^c(\omega')$ for all σ'_i , so $(M, \omega) \models B_i^* \varphi$ iff $(M, \omega') \models B_i^* \varphi$. It immediately follows that $B_i^* \varphi \Rightarrow B_i B_i^* \varphi$ is valid (i.e., true at all states in all structures).

The following abbreviations will be useful in the sequel. Let RAT be an abbreviation for $RAT_1 \wedge \dots \wedge RAT_n$, and let $play(\vec{\sigma})$ be an abbreviation for $play_1(\sigma_1) \wedge \dots \wedge play_n(\sigma_n)$.

2.2 Common Counterfactual Belief of Rationality

We are interested in analyzing strategies being played at states where (1) everyone is rational, (2) everyone counterfactually believes that everyone else is rational (i.e., for every player i , i believes that everyone else would still be rational even if i were to switch strategies), (3) everyone counterfactually believes that everyone else is rational, and counterfactually believes that everyone else is rational, and so on. For each player i , define the formulas $SRAT_i^k$ (player i is strongly k -level rational) inductively, by taking $SRAT_i^0$ to be *true* and $SRAT_i^{k+1}$ to be an abbreviation of

$$RAT_i \wedge B_i^*(\bigwedge_{j \neq i} SRAT_j^k).$$

Let $SRAT^k$ be an abbreviation of $\bigwedge_{j=1}^n SRAT_j^k$.

Define $CCBR$ (common counterfactual belief of rationality) as follows:

- $(M, \omega) \models CCBR$ iff $(M, \omega) \models SRAT^k$ for all $k \geq 1$.

Note that it is critical in the definition of $SRAT_i^k$ that we require only that player i counterfactually believes that everyone else (i.e., the players other than i) are rational, and believe that everyone else is rational, and so on. Player i has no reason to believe that his own strategy would be rational if he were to switch strategies; indeed, $B_i^* RAT_i$ can hold only if *every* strategy for player i is rational with respect to i 's beliefs. This is why we do not define $CCBR$ as $CB^* RAT$.³

We also consider the consequence of just common belief of rationality in our setting. Define $WRAT_i^k$ (player i is weakly k -level rational) just as $SRAT_i^k$, except that B_i^* is replaced by B_i . An easy induction on k shows that $WRAT_i^{k+1}$ implies $WRAT_i^k$ and that $WRAT_i^k$ implies $B_i(WRAT_i^k)$.⁴ It follows that we could have equivalently defined $WRAT_i^{k+1}$ as

$$RAT_i \wedge B_i(\bigwedge_{j=1}^n WRAT_j^k).$$

Thus, $WRAT^{k+1}$ is equivalent to $RAT \wedge EB(WRAT^k)$. As a consequence we have the following:

Proposition 2.1: $(M, \omega) \models CB(RAT)$ iff $(M, \omega) \models WRAT^k$ for all $k \geq 0$.

³Interestingly, Samet [1996] essentially considers an analogue of $CB^* RAT$. This works in his setting since he is considering only events in the past, not events in the future.

⁴We can also show that $SRAT_i^{k+1}$ implies $SRAT_i^k$, but it is not the case that $SRAT_i^k$ implies $B_i^* SRAT_i^k$, since RAT does not imply $B_i^* RAT$.

3 Characterizing Common Counterfactual Belief of Rationality

To put our result into context, we first restate the characterizations of rationalizability given by Tan and Werlang [1988] and Brandenburger and Dekel [1987] in our language. We first recall Pearce’s [1984] definition of rationalizability.

Definition 3.1: A strategy σ_i for player i is *rationalizable* if, for each player j , there is a set $\mathcal{Z}_j \subseteq \Sigma_j(\Gamma)$ and, for each strategy $\sigma'_j \in \mathcal{Z}_j$, a probability measure $\mu_{\sigma'_j}$ on $\Sigma_{-j}(\Gamma)$ whose support is a subset of \mathcal{Z}_{-j} such that

- $\sigma_i \in \mathcal{Z}_i$; and
- for strategy $\sigma'_j \in \mathcal{Z}_j$, strategy σ'_j is a best response to (the beliefs) $\mu_{\sigma'_j}$.

A strategy profile $\vec{\sigma}$ is rationalizable if every strategy σ_i in the profile is rationalizable. ■

Theorem 3.2: [Brandenburger and Dekel 1987; Tan and Werlang 1988] $\vec{\sigma}$ is rationalizable in a game Γ iff there exists a finite probability structure M that is appropriate for Γ and a state ω such that $(M, \omega) \models \text{play}(\vec{\sigma}) \wedge CB(RAT)$.

We now consider counterfactual structures. We here provide a condition on the closest-state function under which common (counterfactual) belief of rationality characterizes rationalizable strategies.

3.1 Counterfactual Structures Respecting Unilateral Deviations

Let $M = (\Omega, f, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$ be a finite counterfactual structure that is appropriate for Γ . M *respects unilateral deviations* if, for every state $\omega \in \Omega$, player i , and strategy σ'_i for player i , $s_{-i}(f(\omega, i, \sigma')) = s_{-i}(\omega)$ and $\mathcal{PR}_{-i}(f(\omega, i, \sigma')) = \mathcal{PR}_{-i}(\omega)$; that is, in the closest state to ω where player i switches strategies, everybody else’s strategy and beliefs remain same.

Recall that \mathbb{L}^0 is defined exactly like \mathbb{L} except that we exclude the “counterfactual” modal operators B^* and CB^* . The following theorem shows that for formulas in \mathbb{L}^0 , counterfactual structures respecting unilateral deviations behave just as (standard) probability structures.

Theorem 3.3: For every $\varphi \in \mathbb{L}^0$, there exists a finite probability structure M appropriate for Γ and a state ω such that $(M, \omega) \models \varphi$ iff there exists a finite counterfactual structure M' (strongly) appropriate for Γ that respects unilateral deviations, and a state ω' such that $(M', \omega') \models \varphi$.

Proof: For the “if” direction, let $M' = (\Omega, f, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$ be a finite counterfactual structure that is counterfactually appropriate for Γ (but not necessarily strongly counterfactually appropriate) and respects unilateral deviations. Define $M = (\Omega, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$. Clearly M is a finite probability structure appropriate for Γ ; it follows by a straightforward induction on the length of φ that $(M', \omega) \models \varphi$ iff $(M, \omega) \models \varphi$.

For the “only-if” direction, let $M = (\Omega, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$ be a finite probability structure, and let $\omega \in \Omega$ be a state such that $(M, \omega) \models \varphi$. We assume without loss of generality that for each strategy profile $\vec{\sigma}'$ there exists some state $\omega_{\vec{\sigma}'} \in \Omega$ such that $s(\omega_{\vec{\sigma}'}) = \vec{\sigma}'$ and for each player i , $\mathcal{PR}_i(\omega_{\vec{\sigma}'}) = 1$. (If such a state does not exist, we can always add it.)

We define a finite counterfactual structure $M' = (\Omega', f', \mathcal{PR}'_1, \dots, \mathcal{PR}'_n)$ as follows:

- $\Omega' = \{(\vec{\sigma}', \omega') : \vec{\sigma}' \in \Sigma(\Gamma), \omega' \in \Omega\}$;
- $s'(\vec{\sigma}', \omega') = \vec{\sigma}'$;
- $f((\vec{\sigma}', \omega'), i, \sigma''_i) = ((\sigma''_i, \sigma'_{-i}), \omega')$
- \mathcal{PR}'_i is defined as follows.
 - $\mathcal{PR}'_i(s(\omega'), \omega')(s(\omega''), \omega'') = \mathcal{PR}_i(\omega')(\omega'')$
 - If $\vec{\sigma}' \neq s(\omega')$, $\mathcal{PR}'_i(\vec{\sigma}', \omega')(\vec{\sigma}', \omega_{\vec{\sigma}'}) = 1$.

It follows by construction that M' is strongly appropriate for Γ and respects unilateral deviations. Furthermore, it follows by an easy induction on the length of the formula φ' that for every state $\omega \in \Omega$, $(M, \omega) \models \varphi'$ iff $(M', (s(\omega), \omega)) \models \varphi'$. ■

We can now use Theorem 3.3 together with the standard characterization of common belief of rationality (Theorem 3.2) to characterize both common belief of rationality and common counterfactual belief of rationality.

Theorem 3.4: *The following are equivalent:*

- (a) $\vec{\sigma}$ is rationalizable in Γ ;
- (b) there exists a finite counterfactual structure M that is appropriate for Γ and respects unilateral deviations, and a state ω such that $(M, \omega) \models \text{play}(\vec{\sigma}) \wedge_{i=1}^n \text{WRAT}_i^k$ for all $k \geq 0$;
- (c) there exists a finite counterfactual structure M that is strongly appropriate for Γ and respects unilateral deviations and a state ω such that $(M, \omega) \models \text{play}(\vec{\sigma}) \wedge_{i=1}^n \text{WRAT}_i^k$ for all $k \geq 0$;

- (d) *there exists a finite counterfactual structure M that is appropriate for Γ and respects unilateral deviations and a state ω such that $(M, \omega) \models \text{play}(\vec{\sigma}) \wedge_{i=1}^n \text{SRAT}_i^k$ for all $k \geq 0$;*
- (e) *there exists a finite counterfactual structure M that is strongly appropriate for Γ and respects unilateral deviations and a state ω such that $(M, \omega) \models \text{play}(\vec{\sigma}) \wedge_{i=1}^n \text{SRAT}_i^k$ for all $k \geq 0$.*

Proof: The equivalence of (a), (b), and (c) is immediate from Theorem 3.2, Theorem 3.3, and Proposition 2.1. We now prove the equivalence of (b) and (d). Consider an counterfactual structure M that is appropriate for Γ and respects unilateral deviations. The result follows immediately once we show that for all states ω and all $i \geq 0$, $(M, \omega) \models \text{WRAT}_i^k$ iff $(M, \omega) \models \text{SRAT}_i^k$. An easy induction on k shows that $\text{SRAT}_i^k \Rightarrow \text{WRAT}_i^k$ is valid in all counterfactual structures, not just ones that respect unilateral deviations. We prove the converse in structures that respect unilateral deviations by induction on k . The base case holds trivially. For the induction step, suppose that $(M, \omega) \models \text{WRAT}_i^k$; that is, $(M, \omega) \models \text{RAT}_i \wedge B_i(\wedge_{j \neq i} \text{WRAT}_j^{k-1})$. Thus, for all $\omega' \in \text{Supp}(\mathcal{PR}_i(\omega))$, we have that $(M, \omega') \models \wedge_{j \neq i} \text{WRAT}_j^{k-1}$. Thus, by the induction hypothesis, $(M, \omega') \models \wedge_{j \neq i} \text{SRAT}_j^{k-1}$. Since, as we have observed, the truth of a formula of the form $B_j^* \varphi$ at a state ω'' depends only on j 's beliefs at ω'' and the truth of RAT_j depends only on j 's strategy and beliefs at ω'' , it easily follows that, if j has the same beliefs and plays the same strategy at ω_1 and ω_2 , then $(M, \omega_1) \models \text{SRAT}_j^{k-1}$ iff $(M, \omega_2) \models \text{SRAT}_j^{k-1}$. Since $(M, \omega') \models \wedge_{j \neq i} \text{SRAT}_j^{k-1}$ and M respect unilateral deviations, for all strategies σ'_i , it follows that $(M, f(\omega', i, \sigma'_i)) \models \wedge_{j \neq i} \text{SRAT}_j^{k-1}$. Thus, $(M, \omega) \models \text{RAT}_i \wedge B_i^*(\wedge_{j \neq i} \text{SRAT}_j^{k-1})$, as desired. The argument that (c) is equivalent to (e) is identical; we just need to consider strongly appropriate counterfactual structures rather than just appropriate counterfactual structures. ■

Remark 3.5: Note that, in the proofs of Theorems 3.3 and 3.4, a weaker condition on the counterfactual structure would suffice, namely, that we restrict to counterfactual structures where, for every state $\omega \in \Omega$, player i , and strategy σ'_i for player i , the projection of $\mathcal{PR}_{i, \sigma'_i}^c(\omega)$ onto strategies and beliefs of players $-i$ is equal to the projection of $\mathcal{PR}_i(\omega)$ onto strategies and beliefs of players $-i$. That is, every player's counterfactual beliefs regarding other players' strategies and beliefs are the same as the player's actual beliefs.

3.2 Iterated Minimax Domination

We now characterize common counterfactual belief of rationality without putting any restrictions on the counterfactual structures (other than them being appropriate, or strongly appropriate). Our characterization is based on ideas that come from the characterization of rationalizability. It is well known that rationalizability can be characterized in terms of an iterated deletion procedure, where at each stage, a strategy σ for player i is deleted if there are no beliefs that i could have about the undeleted strategies for the players other than i that would make σ rational [Pearce 1984]. Thus, there is a deletion procedure that, when applied repeatedly, results in only the rationalizable strategies, that is, the strategies that are played in states where there is common belief of rationality, being left undeleted. We now show that there is an analogous way of characterizing common counterfactual belief of rationality.

The key to our characterization is the notion of *minimax dominated* strategies.

Definition 3.6: Strategy σ_i for player i in game Γ is *minimax dominated* with respect to $\Sigma'_{-i} \subseteq \Sigma_{-i}(\Gamma)$ iff there exists a strategy $\sigma'_i \in \Sigma_i(\Gamma)$ such that

$$\min_{\tau_{-i} \in \Sigma'_{-i}} u_i(\sigma'_i, \tau_{-i}) > \max_{\tau_{-i} \in \Sigma'_{-i}} u_i(\sigma_i, \tau_{-i}).$$

■

In other words, player i 's strategy σ is minimax dominated with respect to Σ'_{-i} iff there exists a strategy σ' such that the worst-case payoff for player i if he uses σ' is strictly better than his best-case payoff if he uses σ , given that the other players are restricted to using a strategy in Σ'_{-i} .

In the standard setting, if a strategy σ_i for player i is dominated by σ'_i then we would expect that a rational player will never play σ_i , because σ'_i is a strictly better choice. As is well known, if σ_i is dominated by σ'_i , then there are no beliefs that i could have regarding the strategies used by the other players according to which σ_i is a best response [Pearce 1984]. This is no longer the case in our setting. For example, in the standard setting, cooperation is dominated by defection in Prisoner's Dilemma. But in our setting, suppose that player 1 believes that if he cooperates, then the other player will cooperate, while if he defects, then the other player will defect. Then cooperation is not dominated by defection.

So when can we guarantee that playing a strategy is irrational in our setting? This is the case only if the strategy is minimax dominated. If σ_i is minimax dominated by σ'_i , there are no counterfactual beliefs that i could have that would justify playing σ_i . Conversely, if σ_i is not minimax dominated by any strategy, then there are beliefs and counterfactual beliefs that i could have that would justify playing

σ_i . Specifically, i could believe that the players in $-i$ are playing the strategy profile that gives i the best possible utility when he plays σ_i , and that if he switches to another strategy σ'_i , the other players will play the strategy profile that gives i the worst possible utility given that he is playing σ'_i .

Note that we consider only domination by pure strategies. It is easy to construct examples of strategies that are not minimax dominated by any pure strategy, but are minimax dominated by a mixed strategy. Our characterization works only if we restrict to domination by pure strategies. The characterization, just as with the characterization of rationalizability, involves iterated deletion, but now we do not delete dominated strategies in the standard sense, but minimax dominated strategies.

Definition 3.7: Define $NSD_j^k(\Gamma)$ inductively: let $NSD_j^0(\Gamma) = \Sigma_j$ and let $NSD_j^{k+1}(\Gamma)$ consist of the strategies in $NSD_j^k(\Gamma)$ not minimax dominated with respect to $NSD_{-j}^k(\Gamma)$. Strategy σ survives k rounds of iterated deletion of minimax dominated strategies for player i if $\sigma \in NSD_i^k(\Gamma)$. Strategy σ for player i survives iterated deletion of minimax dominated strategies if it survives k rounds of iterated deletion of strongly dominated for all k , that is, if $\sigma \in NSD_i^\infty(\Gamma) = \cap_k NSD_i^k(\Gamma)$. ■

In the deletion procedure above, at each step we remove *all* strategies that are minimax dominated; that is we perform a “maximal” deletion at each step. As we now show, the set of strategies that survives iterated deletion is actually independent of the deletion order.

Let S^0, \dots, S^m be sets of strategy profiles. $\vec{S} = (S^0, S^1, \dots, S^m)$ is a *terminating deletion sequence* for Γ if, for $j = 0, \dots, m-1$, $S^{j+1} \subset S^j$ (note that we use \subset to mean proper subset) and all players i , S_i^{j+1} contains all strategies for player i not minimax dominated with respect to S_{-i}^j (but may also contain some strategies that are minimax dominated), and S_i^m does not contain any strategies that are minimax dominated with respect to S_{-i}^m . A set T of strategy profiles has *ambiguous* terminating sets if there exist two terminating deletion sequences $\vec{S} = (T, S_1, \dots, S_m)$, $\vec{S}' = (T, S'_1, \dots, S'_{m'})$ such that $S_m \neq S'_{m'}$; otherwise, we say that T has a *unique terminating set*.

Proposition 3.8: *No (finite) set of strategy profiles has ambiguous terminating sets.*

Proof: Let T be a set of strategy profiles of least cardinality that has ambiguous terminating deletion sequences $\vec{S} = (T, S_1, \dots, S_m)$ and $\vec{S}' = (T, S'_1, \dots, S'_{m'})$, where $S_m \neq S'_{m'}$. Let T' be the set of strategies that are not minimax dominated with respect to T . Clearly $T' \neq \emptyset$ and, by definition, $T' \subseteq S_1 \cap S'_1$. Since T' , S_1 , and S'_1 all have cardinality less than that of T , they must all have unique terminating sets; moreover, the terminating sets must be the same. For consider a

terminating deletion sequence starting at T' . We can get a terminating deletion sequence starting at S_1 by just appending this sequence to S_1 (or taking this sequence itself, if $S_1 = T'$). We can similarly get a terminating deletion sequence starting at S'_1 . Since all these terminating deletion sequences have the same final element, this must be the unique terminating set. But (S_1, \dots, S_m) and $(S'_1, \dots, S'_{m'})$ are terminating deletion sequences with $S_m \neq S'_{m'}$, a contradiction. ■

Corollary 3.9: *The set of strategies that survives iterated deletion of minimax dominated strategies is independent of the deletion order.*

Remark 3.10: Note that in the definition of $NSD_i^k(\Gamma)$, we remove all strategies that are dominated by some strategy in $\Sigma_i(\Gamma)$, not just those dominated by some strategy in $NSD_i^{k-1}(\Gamma)$. Nevertheless, the definition would be equivalent even if we had considered only dominating strategies in $NSD_i^{k-1}(\Gamma)$. For suppose not. Let k be the smallest integer such that there exists some strategy $\sigma_i \in NSD_i^{k-1}(\Gamma)$ that is minimax dominated by a strategy $\sigma'_i \notin NSD_i^{k-1}(\Gamma)$, but there is no strategy in $NSD_i^{k-1}(\Gamma)$ that dominates σ_i . That is, σ'_i was deleted in some previous iteration. Then there exists a sequence of strategies $\sigma_i^0, \dots, \sigma_i^q$ and indices $k_0 < k_1 < \dots < k_q = k - 1$ such that $\sigma_i^0 = \sigma'_i$, $\sigma_i^j \in NSD_i^{k_j}(\Gamma)$, and for all $0 \leq j < q$, σ_i^j is minimax dominated by σ_i^{j+1} with respect to $NSD_i^{k_j-1}(\Gamma)$. Since $NSD_i^{k-2}(\Gamma) \subseteq NSD_i^j(\Gamma)$ for $j \leq k - 2$, an easy induction on j shows that σ_i^q minimax dominates σ_i^{q-j} with respect to NSD_i^{k-2} for all $0 < j \leq q$. In particular, σ_i^q minimax dominates $\sigma_i^0 = \sigma'_i$ with respect to NSD_i^{k-2} . ■

The following example shows that iteration has bite: there exist a 2-player game where each player has k actions and $k - 1$ rounds of iterations are needed.

Example 3.11: Consider a two-player game, where both players announce a value between 1 and k . Both players receive in utility the smallest of the values announced; additionally, the player who announces the larger value get a reward of $p = 0.5$.⁵ That is, $u(x, y) = (y + p, y)$ if $x > y$, $(x, x + p)$ if $y > x$, and (x, x) if $x = y$. In the first step of the deletion process, 1 is deleted for both players; playing 1 can yield a max utility of 1, whereas the minimum utility of any other action is 1.5. Once 1 is deleted, 2 is deleted for both players: 2 can yield a max utility of 2, and the min utility of any other action (once 1 is deleted) is 2.5. Continuing this process, we see that only (k, k) survives. ■

⁵This game can be viewed as a reverse variant of the Traveler's dilemma [Basu 1994], where the player who announces the smaller value gets the reward.

3.3 Characterizing Iterated Minimax Domination

We now show that strategies surviving iterated removal of minimax dominated strategies characterize the set of strategies consistent with common counterfactual belief of rationality in (strongly) appropriate counterfactual structures. As a first step, we define a “minimax” analogue of rationalizability.

Definition 3.12: A strategy profile $\vec{\sigma}$ in game Γ is *minimax rationalizable* if, for each player i , there is a set $\mathcal{Z}_i \subseteq \Sigma_i(\Gamma)$ such that

- $\sigma_i \in \mathcal{Z}_i$;
- for every strategy $\sigma'_i \in \mathcal{Z}_i$ and strategy $\sigma''_i \in \Sigma_i(\Gamma)$,

$$\max_{\tau_{-i} \in \mathcal{Z}_{-i}} u_i(\sigma'_i, \tau_{-i}) \geq \min_{\tau_{-i} \in \mathcal{Z}_{-i}} u_i(\sigma''_i, \tau_{-i}).$$

■

Theorem 3.13: *The following are equivalent:*

- (a) $\vec{\sigma} \in NSD^\infty(\Gamma)$;
- (b) $\vec{\sigma}$ is minimax rationalizable in Γ ;
- (c) there exists a finite counterfactual structure M that is strongly appropriate for Γ and a state ω such that $(M, \omega) \models \text{play}(\vec{\sigma}) \wedge_{i=1}^n \text{SRAT}_i^k$ for all $k \geq 0$;
- (d) for all players i , there exists a finite counterfactual structure M that is appropriate for Γ and a state ω such that $(M, \omega) \models \text{play}_i(\sigma_i) \wedge \text{SRAT}_i^k$ for all $k \geq 0$.

Proof: We prove that (a) implies (b) implies (c) implies (d) implies (a). We first introduce some helpful notation. Recall that $\arg \max_x f(x) = \{y : \text{for all } z, f(z) \leq f(y)\}$; $\arg \min_x f(x)$ is defined similarly. For us, x ranges over pure strategies or pure strategy profiles, and we will typically be interested in considering some element of the set, rather than the whole set. Which element we take does not matter. We thus assume that there is some order on the set of pure strategies and strategy profiles, and take the $\arg \max_x^* f(x)$ to be the maximum element of $\arg \max_x f(x)$ with respect to this order; $\arg \min_x^* f(x)$ is defined similarly.

(a) \Rightarrow (b): Let K be an integer such that $NSD^K(\Gamma) = NSD^{K+1}(\Gamma)$; such a K must exist since the game is finite. It also easily follows that for each player j , $NSD_j^K(\Gamma)$ is non-empty: in iteration $k + 1$, no NSD_j^k -maximin strategy, that is, no strategy in $\arg \max_{\sigma'_j \in NSD_j^k(\Gamma)} \min_{\tau_{-j} \in NSD_j^k(\Gamma)} u_j(\sigma'_j, \tau_{-j})$, is deleted, since no maximin strategy is minimax dominated by a strategy in $NSD_j^k(\Gamma)$ (recall that by Remark 3.10, it suffices to consider domination by strategies in $NSD_j^k(\Gamma)$). Let $\mathcal{Z}'_j = NSD_j^K(\Gamma)$. It immediately follows that the sets $\mathcal{Z}'_1, \dots, \mathcal{Z}'_n$ satisfy the conditions of Definition 3.12.

(b) \Rightarrow (c): Suppose that $\vec{\sigma}$ is minimax rationalizable. Let $\mathcal{Z}_1, \dots, \mathcal{Z}_n$ be the sets guaranteed to exist by Definition 3.12. Let $\mathcal{W}^i = \{(\vec{\sigma}, i) \mid \vec{\sigma} \in \mathcal{Z}_{-i} \times \Sigma_i\}$, and let $\mathcal{W}^0 = \{(\vec{\sigma}, 0) \mid \vec{\sigma} \in \mathcal{Z}_1 \times \dots \times \mathcal{Z}_n\}$. Think of \mathcal{W}^0 as states where everyone is (higher-level) rational, and of \mathcal{W}^i as “counterfactual” states where player i has changed strategies. In states in \mathcal{W}^0 , each player j assigns probability 1 to the other players choosing actions that *maximize* j ’s utility (given his action). On the other hand, in states in \mathcal{W}^i , where $i \neq 0$, player i assigns probability 1 to the other players choosing actions that *minimize* i ’s utility, whereas all other player $j \neq i$ still assign probability 1 to other players choosing actions that maximize j ’s utility.

Define a structure $M = (\Omega, f, s, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$, where

- $\Omega = \cup_{i \in \{0, 1, \dots, n\}} \mathcal{W}^i$;
- $s(\vec{\sigma}', i) = \vec{\sigma}'$;
- $\mathcal{PR}_j(\vec{\sigma}', i)(\vec{\sigma}'', i') = \begin{cases} 1 & \text{if } i = j = i', \sigma'_i = \sigma''_i, \text{ and } \sigma''_{-i} = \arg \min_{\tau_{-i} \in \mathcal{Z}_{-i}}^* u_j(\sigma'_i, \tau_{-i}), \\ 1 & \text{if } i \neq j, i' = 0, \text{ and } \sigma'_j = \sigma''_j, \text{ and } \sigma''_{-j} = \arg \max_{\tau_{-j} \in \mathcal{Z}_{-j}}^* u_j(\sigma'_j, \tau_{-j}), \\ 0 & \text{otherwise;} \end{cases}$
- $f((\vec{\sigma}', i), j, \sigma''_j) = \begin{cases} (\vec{\sigma}', i) & \text{if } \sigma'_j = \sigma''_j, \\ ((\sigma''_j, \tau'_{-j}), j) & \text{otherwise, where } \tau'_{-j} = \arg \min_{\tau_{-j} \in \mathcal{Z}_{-j}}^* u_j(\sigma'_j, \tau_{-j}). \end{cases}$

It follows by inspection that M is strongly appropriate for Γ . We now prove by induction on k that, for all $k \geq 1$ all $i \in \{0, 1, \dots, n\}$, and all states $\omega \in \mathcal{W}^i$, $(M, \omega) \models \bigwedge_{j \neq i} SRAT_j^k$.

For the base case ($k = 1$), since $SRAT_j^1$ is logically equivalent to RAT_j , we must show that if $\omega \in \mathcal{W}^i$, then $(M, \omega) \models \bigwedge_{j \neq i} RAT_j$. Suppose that $\omega = (\vec{\sigma}', i) \in \mathcal{W}^i$. If $i \neq j$, then at ω , player j places probability 1 on the true state being $\omega' = (\vec{\sigma}'', 0)$, where $\sigma''_j = \sigma'_j$ and $\sigma''_{-j} = \arg \max_{\tau_{-j} \in \mathcal{Z}_{-j}}^* u_j(\sigma'_j, \tau_{-j})$. Player j must be rational, since if there exists some strategy τ'_j such that $u_j(\vec{\sigma}'') < \sum_{\omega' \in \Omega} \mathcal{PR}_{j, \tau'_j}^c(\omega)(\omega') u_j(\tau'_j, s_{-j}(\omega'))$, then the definition of \mathcal{PR}_j guarantees that

$u_j(\bar{\sigma}'') < u_j(\tau'_j, \bar{\tau}''_{-j})$, where $\tau''_j = \arg \min_{\tau_{-j} \in \mathcal{Z}_{-j}} u_j(\sigma'_j, \tau_{-j})$. If this inequality held, then τ'_j would minimax dominate σ'_j , contradicting the assumption that $\sigma'_j \in \mathcal{Z}_j$. For the induction step, suppose that the result holds for k ; we show that it holds for $k+1$. Suppose that $\omega \in \mathcal{W}^0$ and $j \neq i$. By construction, the support of $\mathcal{PR}_j(\omega)$ is a subset of \mathcal{W}^0 ; by the induction hypothesis, it follows that $(M, \omega) \models B_j(\bigwedge_{j'=1}^n SRAT_{j'}^k)$. Moreover, by construction, it follows that for all players j and all strategies $\sigma'_j \neq s_i(\omega)$, the support of $\mathcal{PR}_{j, \sigma'_j}^c(\omega)$ is a subset of \mathcal{W}^j . By the induction hypothesis, it follows that for all $j \neq i$, $(M, \omega) \models B_j^*(\bigwedge_{j' \neq j} SRAT_{j'}^k)$. Finally, it follows from the induction hypothesis that for all $j \neq i$, $(M, \omega) \models SRAT_j^k$. Since $SRAT_j^k$ implies RAT_j , it follows that for all $j \neq i$, $(M, \omega) \models RAT_j \wedge B_j^*(\bigwedge_{j' \neq j} SRAT_{j'}^k)$, which proves the induction step.

(c) \Rightarrow (d): The implication is trivial.

(d) \Rightarrow (a): We prove an even stronger statement: For all $k \geq 0$, if there exists a finite counterfactual structure M^k that is appropriate for Γ and a state ω such that $(M^k, \omega) \models play_i(\sigma_i) \wedge SRAT_i^k$, then $\sigma_i \in NSD_i^k(\Gamma)$.⁶ We proceed by induction on k . The result clearly holds if $k = 0$. Suppose that the result holds for $k-1$ for $k \geq 1$; we show that it holds for k . Let $M^k = (\Omega, f, s, \mathcal{P}_1, \dots, \mathcal{P}_n)$ be a finite counterfactual structure that is appropriate for Γ and a state ω such that $(M^k, \omega') \models play_i(\sigma_i) \wedge SRAT_i^k$. Replacing $SRAT_i^k$ by its definition, we get that

$$(M^k, \omega') \models play_i(\sigma_i) \wedge RAT_i \wedge B_i^*(\bigwedge_{j \neq i} SRAT_j^{k-1}).$$

By definition of B_i^* , it follows that for all strategies σ'_i for player i and all ω'' such that $\mathcal{PR}_{i, \sigma'_i}^c(\omega')(\omega'') > 0$,

$$(M^k, \omega'') \models \bigwedge_{j \neq i} SRAT_j^{k-1},$$

so by the induction hypothesis, it follows that for all ω'' such that $\mathcal{PR}_{i, \sigma'_i}^c(\omega')(\omega'') > 0$, we have $s_{-i}(\omega'') \in NSD_{-i}^{k-1}(\Gamma)$. Since $(M^k, \omega') \models play_i(\sigma_i) \wedge RAT_i$, it follows that σ_i cannot be minimax dominated with respect to $NSD_{-i}^{k-1}(\Gamma)$. Since, for all $j' > 1$, $NSD_{-i}^{j'}(\Gamma) \subseteq NSD_{-i}^{j'-1}(\Gamma)$, it follows that, for all $k' < k$, σ_i is not minimax dominated with respect to $NSD_{-i}^{k'}(\Gamma)$. Thus, $\sigma_i \in NSD_i^k(\Gamma)$. ■

⁶The converse also holds; we omit the details.

4 Characterizing Analogues of Nash Equilibrium

In this section, we consider analogues of Nash equilibrium in our setting. This allows us to relate our approach to the work of Tennenholtz [2004] and Kalai et al. [2010]. In the standard setting, if a strategy profile $\vec{\sigma}$ is a Nash equilibrium, then there exists a state where $\vec{\sigma}$ is played, common belief of rationality holds, and additionally, the strategy profile is (commonly) known to the players. To study analogues of Nash equilibrium, we thus investigate the effect of adding assumptions about knowledge of the players' strategies.

We consider several ways of formalizing this. The weakest approach is to simply require that the actual strategies used by the players is known.

- $(M, \omega) \models \text{KS}$ iff, for all players i ,

$$\mathcal{PR}_i(\omega)(\llbracket \mathbf{s}_{-i}(\omega) \rrbracket_M) = 1.$$

KS does not require that player i knows how players $-i$ will respond to i switching strategies. A stronger condition would be to require not only that every player i knows the strategies of the other players, but also how they respond to i switching strategies.

- $(M, \omega) \models \text{KR}$ iff, for all players i and strategies σ'_i for i ,

$$\mathcal{PR}_{i, \sigma'_i}^c(\omega)(\llbracket \mathbf{s}_{-i}(f(\omega, i, \sigma'_i)) \rrbracket_M) = 1.$$

Clearly, KR implies KS (by simply considering $\sigma'_i = \mathbf{s}_i(\omega)$). An even stronger condition is to require that the players know the true state of the world.

- $(M, \omega) \models \text{KW}$ iff, for all players i ,

$$\mathcal{PR}_i(\omega)(\omega) = 1.$$

Note that if all players know the true state of the world, then they also counterfactually know the true state of the world: for every player i and every strategy σ'_i for player i ,

$$\mathcal{PR}_{i, \sigma'_i}^c(\omega)(f(\omega, i, \sigma'_i)) = 1.$$

It follows that KW implies KR and thus also KS. Additionally, note that KW implies $EB(\text{KW})$, so KW also implies $CB(\text{KW})$.

We now characterize $CCBR$ in structures satisfying the conditions above. We say that a strategy profile $\vec{\sigma}$ is *individually rational* (IR) if for every player i in the game Γ ,

$$u_i(\vec{\sigma}) \geq \max_{\sigma'_i \in \Sigma_i(\Gamma)} \min_{\tau_{-i} \in \Sigma_{-i}(\Gamma)} u_i(\sigma'_i, \tau_{-i}).$$

Although every IR strategy profile is contained in $NSD^1(\Gamma)$, it is not necessarily contained in $NSD^2(\Gamma)$. That is, IR strategies may not survive two rounds of deletion of minimax dominated strategies. To see this, consider the game Γ in Example 3.11. Both players' maximin payoff is 1.5, so every strategy profile in $NSD^1(\Gamma) = \{(x, y) \mid 2 \leq x, y \leq k\}$ is IR, but $NSD^2(\Gamma)$ does not contain $(2, 2)$.

As the following simple example shows, not every strategy profile that survives deletion iterated deletion of minimax dominated strategies is IR.

Example 4.1: Consider the game with payoffs given in the table below.

	c	d
a	(100, 0)	(100, 0)
b	(150, 0)	(50, 0)

All strategy profiles survive iterated deletion of minimax dominated strategies, but (b, d) is not individually rational since playing a always guarantees the row player utility 100. ■

Let $IR(\Gamma)$ denote the set of IR strategy profiles in Γ , and let $IR(\mathcal{Z}_1 \times \dots \times \mathcal{Z}_n, \Gamma) = IR(\Gamma')$ where Γ' is the subgame of Γ obtained by restricting player i 's strategy set to \mathcal{Z}_i . That is, $IR(\mathcal{Z}_1 \times \dots \times \mathcal{Z}_n, \Gamma)$ is the set of strategies $\vec{\sigma} \in \mathcal{Z}_1 \times \dots \times \mathcal{Z}_n$ such that for every player i ,

$$u_i(\vec{\sigma}) \geq \max_{\sigma'_i \in \mathcal{Z}_i} \min_{\tau_{-i} \in \mathcal{Z}_{-i}} u_i(\sigma'_i, \tau_{-i}).$$

A stronger way of capturing individual rationality of subgames is to require that the condition above hold even if the max is taken over every $\sigma'_i \in \Sigma(\Gamma)$ (as opposed to only $\sigma'_i \in \mathcal{Z}_i$). More precisely, let $IR'(\mathcal{Z}_1 \times \dots \times \mathcal{Z}_n, \Gamma)$ be the set of strategies $\vec{\sigma} \in \mathcal{Z}_1 \times \dots \times \mathcal{Z}_n$ such that, for all players i ,

$$u_i(\vec{\sigma}) \geq \max_{\sigma'_i \in \Sigma(\Gamma)} \min_{\tau_{-i} \in \mathcal{Z}_{-i}} u_i(\sigma'_i, \tau_{-i}).$$

Our characterization of $CCBR$ in the presence of (common) knowledge of strategies follows.

Theorem 4.2: *The following are equivalent:*

- (a) $\vec{\sigma} \in IR(NSD^\infty(\Gamma), \Gamma)$;
- (b) $\vec{\sigma} \in IR'(NSD^\infty(\Gamma), \Gamma)$;

- (c) $\vec{\sigma}$ is minimax rationalizable and $\vec{\sigma} \in IR'(\mathcal{Z}_1 \times \dots \times \mathcal{Z}_n, \Gamma)$, where $\mathcal{Z}_1, \dots, \mathcal{Z}_n$ are the sets of strategies guaranteed to exist by the definition of minimax rationalizability;
- (d) there exists a finite counterfactual structure M that is strongly appropriate for Γ and a state ω such that $(M, \omega) \models KW \wedge play(\vec{\sigma}) \wedge_{i=1}^n SRAT_i^k$ for every $k \geq 0$;
- (e) there exists a finite counterfactual structure M that is appropriate for Γ and a state ω such that $(M, \omega) \models KS \wedge play(\vec{\sigma}) \wedge_{i=1}^n SRAT_i^k$ for every $k \geq 0$.

Proof: Again, we prove that (a) implies (b) implies (c) implies (d) implies (e) implies (a).

(a) \Rightarrow (b): We show that if $\vec{\sigma} \in IR(NSD^k(\Gamma), \Gamma)$ then $\vec{\sigma} \in IR'(NSD^k(\Gamma), \Gamma)$. The implication then follows from the fact that since the game is finite there exists some K such that $NSD^K(\Gamma) = NSD^\infty(\Gamma)$.

Assume by way of contradiction that $\vec{\sigma} \in IR(NSD^k(\Gamma), \Gamma)$ but $\vec{\sigma} \notin IR'(NSD^k(\Gamma), \Gamma)$; that is, there exists a player i and a strategy $\sigma'_i \notin NSD_i^k(\Gamma)$ such that

$$\min_{\tau_{-i} \in NSD_{-i}^k(\Gamma)} u_i(\sigma'_i, \tau_{-i}) > u_i(\vec{\sigma}).$$

By the argument in Remark 3.10, there exists a strategy $\sigma''_i \in NSD_i^k(\Gamma)$ such that $u_i(\sigma''_i, \tau''_{-i}) > u_i(\sigma'_i, \tau'_{-i})$ for all $\tau''_{-i}, \tau'_{-i} \in NSD_{-i}^k(\Gamma)$. It follows that

$$\min_{\tau_{-i} \in NSD_{-i}^k(\Gamma)} u_i(\sigma''_i, \tau_{-i}) > u_i(\vec{\sigma}).$$

Thus, $\vec{\sigma} \notin IR(NSD^k(\Gamma), \Gamma)$.

(b) \Rightarrow (c): The implication follows in exactly the same way as in the proof that (a) implies (b) in Theorem 3.13.

(c) \Rightarrow (d): Suppose that $\vec{\sigma}$ is minimax rationalizable. Let $\mathcal{Z}_1, \dots, \mathcal{Z}_n$ be the sets guaranteed to exist by Definition 3.12, and suppose that $\vec{\sigma} \in IR'(\mathcal{Z}_1 \times \mathcal{Z}_n, \Gamma)$. Define the sets \mathcal{W}^i as in the proof of Theorem 3.13. Define the structure M just as in the proof of Theorem 3.13, except that for all players i , let $\mathcal{PR}_i((\vec{\sigma}, 0))((\vec{\sigma}', i')) = 1$ in case $\vec{\sigma}' = \vec{\sigma}$ and $i' = 0$. Clearly $(M, (\vec{\sigma}, 0)) \models KW$. It follows using the same arguments as in the proof of Theorem 3.13 that M is strongly appropriate and that $(M, (\vec{\sigma}, 0)) \models play(\vec{\sigma}) \wedge_{i=1}^n SRAT_i^k$ for every $k \geq 0$; we just need to rely on the (strong) IR property of $\vec{\sigma}$ to prove the base case of the induction.

(d) \Rightarrow (e): The implication is trivial.

(e) \Rightarrow (a): Recall that since the game is finite, there exists a constant K such that $NSD^{K-1}(\Gamma) = NSD^K(\Gamma) = NSD^\infty(\Gamma)$. We show that if there exists a finite counterfactual structure M that is appropriate for Γ and a state ω such that $(M, \omega) \models KS \wedge play(\vec{\sigma}) \wedge_{i=1}^n SRAT_i^K$, then $\vec{\sigma} \in IR(NSD^K(\Gamma), \Gamma)$.

Consider some state ω such that $(M, \omega) \models KS \wedge play(\vec{\sigma}) \wedge_{i=1}^n SRAT_i^K$. By Theorem 3.13, it follows that $\vec{\sigma} \in NSD^K(\Gamma)$. For each player i , it additionally follows that $(M, \omega) \models play(\vec{\sigma}) \wedge EB(play(\vec{\sigma})) \wedge RAT_i \wedge B_i^*(\wedge_{j \neq i} SRAT_j^{K-1})$. By Theorem 3.13, it follows that for every strategy σ'_i for i , the support of the projection of $\mathcal{PR}_{i, \sigma'_i}^c(\omega)$ onto strategies for players $-i$ is a subset of $NSD_{-i}^{K-1}(\Gamma) = NSD_{-i}^K(\Gamma)$. Thus, we have that for every σ'_i , there exists $\tau_{-i} \in NSD_{-i}^K(\Gamma)$ such that $u_i(\vec{\sigma}) \geq u_i(\sigma'_i, \tau_{-i})$, which means that $\vec{\sigma}$ is IR in the subgame induced by restricting the strategy set to $NSD^K(\Gamma)$. ■

It is worth comparing Theorem 4.2 to the results of Tennenholtz [2004] and Kalai et al. [2010] on program equilibria/equilibria with conditional commitments. Recall that these papers focus on 2-player games. In Tennenholtz's model, each player i deterministically picks a program Π_i ; player i 's action is $\Pi_i(\Pi_{-i})$. In the two-player case, a program equilibrium is a pair of programs (Π_1, Π_2) such that no player can improve its utility by unilaterally changing its program. In this model any IR strategy profile (a_1, a_2) can be sustained in a program equilibrium: each player uses the program Π , where $\Pi(\Pi')$ outputs a_i if $\Pi' = \Pi$, and otherwise "punishes" the other player using his minmax strategy. (Tennenholtz extends this result to show that any mixed IR strategy profile can be sustained in a program equilibrium, by considering randomizing programs; Kalai et al. show that all correlated IR strategy profiles can be sustained, by allowing the players to pick a distribution over programs.) In contrast, in our model, a smaller set of strategy profiles can be sustained. This difference can be explained as follows. In the program equilibrium model a player may "punish" the other player using an arbitrary action (e.g., using minimax punishment) although this may be detrimental for him. Common counterfactual belief of rationality disallows such punishments. More precisely, it allows a player i to punish other players only by using a strategy that is rational for player i . On the other hand, as we now show, if we require only common belief (as opposed to counterfactual belief) in rationality, then any IR strategy can be sustained in an equilibrium in our model.

Theorem 4.3: *The following are equivalent:*

- (a) $\vec{\sigma} \in IR(\Gamma)$;

- (b) *there exists a finite counterfactual structure M that is strongly appropriate for Γ and a state ω such that $(M, \omega) \models \text{KW} \wedge \text{play}(\vec{\sigma}) \wedge \text{CB}(\text{RAT})$;*
- (c) *there exists a finite counterfactual structure M that is appropriate for Γ and a state ω such that $(M, \omega) \models \text{KS} \wedge \text{play}(\vec{\sigma}) \wedge \text{CB}(\text{RAT})$.*

Proof: Again, we prove that (a) implies (b) implies (c) implies (a).

(a) \Rightarrow (b): Define a structure $M = (\Omega, f, s, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$, where

- $\Omega = \Sigma(\Gamma)$;
- $s(\vec{\sigma}') = \vec{\sigma}'$;
- $\mathcal{PR}_j(\vec{\sigma}')(\vec{\sigma}') = 1$.
- $f(\vec{\sigma}', i, \sigma_j'') = \begin{cases} \vec{\sigma}' & \text{if } \sigma_j' = \sigma_j'', \\ (\sigma_j'', \tau_{-j}') & \text{otherwise, where } \tau_{-j}' = \arg \min_{\tau_{-j} \in \Sigma_{-j}(\Gamma)} u_j(\sigma_j', \tau_{-j}). \end{cases}$

It follows that M is strongly appropriate for Γ and that $(M, \vec{\sigma}) \models \text{KW}$. Moreover, $(M, \vec{\sigma}) \models \text{RAT}$ since $\vec{\sigma}$ is individually rational; furthermore, since each player considers only the state $\vec{\sigma}$ possible at $\vec{\sigma}$, it follows that $(M, \vec{\sigma}) \models \text{CB}(\text{RAT})$.

(b) \Rightarrow (c): The implication is trivial.

(c) \Rightarrow (a): Suppose that $M = (\Omega, f, s, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$ is a finite counterfactual structure appropriate for Γ , and $(M, \omega) \models \text{KW} \wedge \text{play}(\vec{\sigma}) \wedge \text{CB}(\text{RAT})$. It follows that for each player i , $(M, \omega) \models \text{play}(\vec{\sigma}) \wedge \text{EB}(\text{play}(\vec{\sigma})) \wedge \text{RAT}_i$. Thus, we have that for all strategies σ_i' , there exists $\tau_{-i} \in \Sigma_{-i}(\Gamma)$ such that $u_i(\vec{\sigma}) \geq u_i(\sigma_i', \tau_{-i})$, which means that $\vec{\sigma}$ is IR. ■

5 Discussion

We have introduced a game-theoretic framework for analyzing scenarios where a player may believe that if he were to switch strategies, this intention to switch may be detected by the other players, resulting in them also switching strategies. Our formal model allows players' counterfactual beliefs (i.e., their beliefs about the state of the world in the event that they switch strategies) to be arbitrary—they may be completely different from the players' actual beliefs.

We may also consider a more restricted model where we require that a player i 's counterfactual beliefs regarding other players' strategies and beliefs is ϵ -close to player i 's actual beliefs in total variation distance⁷—that is, for every state $\omega \in \Omega$, player i , and strategy σ'_i for player i , the projection of $\mathcal{PR}_{i,\sigma'_i}^c(\omega)$ onto strategies and beliefs of players $-i$ is ϵ -close to the projection of $\mathcal{PR}_i(\omega)$ onto strategies and beliefs of players $-i$.

We refer to counterfactual structures satisfying this property as ϵ -counterfactual structures. Roughly speaking, ϵ -counterfactual structures restrict to scenarios where players are not “too” transparent to one another; this captures the situation when a player assigns only some “small” probability to its switch in action being noticed by the other players. 0-counterfactual structures behave just as counterfactual structures that respect unilateral deviations: common counterfactual belief of rationality in 0-counterfactual structures characterizes rationalizable strategies (see Remark 3.5). The general counterfactual structures investigated in this paper are 1-counterfactual structures (that is, we do not impose any conditions on players' counterfactual beliefs). We remark that although our characterization results rely on the fact that we consider 1-counterfactual structures, the motivating example in the introduction (the translucent prisoner's dilemma game) shows that even considering ϵ -counterfactual structures with a small ϵ can result in there being strategies consistent with common counterfactual belief of rationality that are not rationalizable. We leave an exploration of common counterfactual belief of rationality in ϵ -counterfactual structures for future work.

References

- Aumann, R. J. (1995). Backwards induction and common knowledge of rationality. *Games and Economic Behavior* 8, 6–19.
- Basu, K. (1994). The traveler's dilemma: paradoxes of rationality in game theory. *American Economic Review* 84(2), 391–395.
- Brandenburger, A. and E. Dekel (1987). Rationalizability and correlated equilibria. *Econometrica* 55, 1391–1402.
- Ekman, P. and W. Friesen (1969). Nonverbal leakage and clues to deception. *Psychiatry* 32, 88–105.
- Halpern, J. Y. (1999). Hypothetical knowledge and counterfactual reasoning. *International Journal of Game Theory* 28(3), 315–330.

⁷Recall that two probability distribution are ϵ -close in total variation distance if the probabilities that they assign to any event E differ by at most ϵ .

- Howard, N. (1971). *Paradoxes of Rationality: Theory of Metagames and Political Behavior*. Cambridge, Mass.: MIT Press.
- Kalai, A., E. Kalai, E. Lehrer, and D. Samet (2010). A commitment folk theorem. *Games and Economic Behavior* 69(1), 127–137.
- Lewis, D. K. (1973). *Counterfactuals*. Cambridge, Mass.: Harvard University Press.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52(4), 1029–1050.
- Salcedo, B. (2013). Implementation without commitment in moral hazard environments. Working paper.
- Samet, D. (1996). Hypothetical knowledge and games with perfect information. *Games and Economic Behavior* 17, 230–251.
- Spohn, W. (2003). Dependency equilibria and the causal structure of decision and game situations. In *Homo Oeconomicus XX*, pp. 195–255.
- Stalnaker, R. C. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in Logical Theory*, pp. 98–112. Blackwell.
- Stalnaker, R. C. (1996). Knowledge, belief and counterfactual reasoning in games. *Economics and Philosophy* 12, 133–163.
- Tan, T. and S. Werlang (1988). The Bayesian foundation of solution concepts of games. *Journal of Economic Theory* 45(45), 370–391.
- Tennenholz, M. (2004). Program equilibrium. *Games and Economic Behavior* 49(12), 363–373.
- von Neumann, J. and O. Morgenstern (1947). *Theory of Games and Economic Behavior* (2nd ed.). Princeton, N.J.: Princeton University Press.
- Zambrano, E. (2004). Counterfactual reasoning and common knowledge of rationality in normal form. *Topics in Theoretical Economics* 4(1).