Submitted to *Theoretical Economics*

Dependency Equilibria: Extending Nash Equilibria to Entangled Belief Systems

WOLFGANG SPOHN
Department of Philosophy, University of Konstanz

MANTAS RADZVILAS
Department of Philosophy, University of Konstanz

GERARD ROTHFUS
Philosophy, Politics, and Economics Program, UNC, Chapel Hill

This paper gives a detailed formal characterization of dependency equilibria — a novel solution concept for games that provides a natural extension of Nash equilibria to strategic interactions where the standard assumption of non-cooperative game theory of the causal independence of the players' choices is retained, but the assumption of their probabilistic independence is forgone. Hence, players' beliefs may be entangled, i.e., permit probabilistic dependencies between their choices, in which case they maximize conditional expected utility (in contrast to correlated equilibria, where players maximize posterior unconditional expected utility). We demonstrate how this novel equilibrium concept can account for seeming out-of-equilibrium behavior in a variety of experimentally and socially relevant games. We further provide lower and upper bounds for the existence of dependency equilibria, determine epistemic conditions for their obtaining, and demonstrate how certain simple iterative belief revision algorithms can lead players into a common dependency equilibrium state.

KEYWORDS. dependency equilibria, Nash equilibria, prisoners' dilemma, epistemic game theory, evolutionary dynamics, deliberational dynamics.

1. INTRODUCTION

Nash equilibria (NE) are the cornerstone of non-cooperative game theory. Characteristically, they assume the causal independence of players' strategy choices. This is reflected in the probabilistic independence of players' mixed strategies in a NE. Or, if one interprets a NE as an equilibrium of beliefs, as in epistemic game theory, causal independence is taken to imply that each player views the choices of all players as evidentially independent of one another. This has been understood since NE were first conceived.

Wolfgang Spohn: wolfgang.spohn@uni-konstanz.de
Mantas Radzvilas: mantas.radzvilas@uni-konstanz.de
Gerard Rothfus: gjrothfu@unc.edu

We gratefully acknowledge funding of this research by the German Research Foundation, Grant No. SP 279/21–1 (Project No. 420094936).

Indeed, the indicated step from the causal to the probabilistic independence of mixed strategies has been taken for granted as well and has rarely been felt to require justification.

In fact, though, this probabilistic independence is a strong and unrealistic assumption that is likely to be violated. We should instead envisage players as having entangled belief systems that allow for probabilistic dependencies amongst their actions or strategy choices.[1] The choices of other players are then treated, in decision theoretic terms, not as independent, but as dependent states of the world. How should an agent rationally behave given such entangled belief systems? By maximizing conditional expected utility. This is the general decision rule applicable when states are act-dependent (see Fishburn (1964)). It reduces to maximizing unconditional expected utility only when the expectation is taken with respect to probabilistically independent states of the world, as developed by Savage (1954). In a NE, each player attempts to maximize his or her unconditional expected utility. However, in entangled belief systems, where players view their choices as probabilistically dependent, only the conditional sense of maximizing expected utility applies. In a dependency equilibrium (DE), as proposed here, each player attempts to maximize conditional expected utility and can do no better in this regard by changing to another choice. So, DE are entirely in the spirit of best-response reasoning, preserving the core idea of NE and extending their application to cases in which players' choices are probabilistically dependent. It is this extension we study in this paper.

One may interject that Aumann (1974) has already provided an equilibrium notion designed for probability distributions in which the choices of players are not independent, namely the by now standard notion of a correlated equilibrium (CE). We will carefully explain the distinction between CE and DE in Section 4.1 and argue that the treatment of dependence should not be left to CE, but is significantly enhanced by DE.

Why should we consider entangled belief systems and the novel equilibrium notion they entail? It seems that the independence in NE is an extreme limiting case, just as would be full dependence. Reality mostly moves between the extremes. It should be clear that entangled belief systems are ubiquitous. In the comparative Section 3 we shall discuss a variety of phenomena displaying dependence and proposals trying to cope with it. We will find that they either are more complicated or fail to be general. They all differ from our treatment.

Alas, these phenomena and considerations are not well reputed in standard game theory. Why? This is because the causal independence of players' strategies is constitutive of non-cooperative games in normal form. And it is taken to entail the probabilistic independence of these strategies. Thus, entangled belief systems as envisaged here appear to exemplify something like magical thinking, the false notion that a player's action can causally influence her opponents' actions.

However, inferring probabilistic from causal independence is generally a fallacy. According to the widely accepted common cause principle of Reichenbach (1956) (ch. 19) every probabilistic correlation must be either a causal one itself or generated by some

[1] When speaking of entangled belief systems, we mean just this. No allusion to quantum theoretic entanglement is intended. Quantum game theory is foreign to this paper.

common cause.[2] So, if the entangled beliefs are to have a causal base, it need not come from the players' belief that their actions causally influence one another's actions, which would contradict the presupposition of games in normal form; it may rather come from the players' belief that they are mutually influenced by a common cause.

Still, it seems that this inference can be saved in the case at hand. One may point out that the players' actions or choices are exogenous variables in decision or game models. They are, in the terminology of causal Bayes net theory, interventions that are exogenous (= parent-less) in the relevant causal Bayes net since they are truncated from all causal predecessors they might have had.[3] This is the distinctive characteristic of causal decision theory to which game theorists are attached. Hence, in the players' eyes, there can't be a common cause of their choices. This entails that the exogenous choices can be evidentially relevant only for their causal effects and excludes any probabilistic dependence between the players' actions. Hence, inferring probabilistic from causal independence seems justified in the case of NE.

Alas, this argument is faulty. Spohn (2003) has suggested that there still may be a common cause of the players' choices. Individual decision theory must conceive the action or choice of an agent as caused by her mental set-up, her beliefs and desires or probabilities and utilities, i.e., precisely by her subjective model of her decision situation. Game theory may conceive of players' choices as being caused in the same way. Prior to their choices, there is plenty of time for causal interaction between players' mental set-ups. If so, these interacting mental set-ups provide a common cause of their choices. For instance, in the original PD story the two criminals have a long career of joint crimes and plenty of occasions to form entangled beliefs which they need not give up when separated by the police in their prison cells. In some such way, the probabilistic dependence allowed in dependency equilibria may have a causal background after all. This is the basic reason why entangled belief systems and their associated DE should be taken seriously even within the setting of non-cooperative game theory and its causal presuppositions. We shall unfold this point in Section 4.1.

This paper begins developing the theory of DE in various directions. It will proceed as follows: Section 2 presents basic definitions, theorems, and examples. In Section 2.1, we define and compare NE, CE and DE in the case of completely mixed strategies; in this case the comparison is particularly palpable. In Section 2.2, we remove the restriction from the previous section and give two equivalent general definitions of DE, one in terms of Popper measures and one in terms of lexicographic probabilities. Moreover, we specify tightest lower and upper bounds for the existence of DE (which, however, will leave a large space in between). Section 2.3 looks at three prominent two-person games in order to showcase the novel DE solution concept in action. It behaves in significantly different ways; see in particular our discussion of PD and the ultimatum game.

Section 3 offers two kinds of rationalization of DE. In Section 3.1, we copy epistemic game theory by showing that common or, rather, second-order mutual knowledge of

[2]Quantum theoretical phenomena like the Einstein-Podolsky-Rosen paradox seem to be exceptions. There is some debate in philosophy whether there are also 'ordinary' exceptions (see Cartwright (2007)), but this need not concern us here.

[3]This is the core of the prevailing interventionist theory of causation, as paradigmatically developed, e.g., by Pearl (2009). See there in particular Section 3.2.3.

the game, of rationality (in the form of maximizing conditional expected utility), and of entangled beliefs entails a DE. This rationalization works just as well as in the case of NE. It may be criticized for its strong common knowledge assumptions. Therefore, the much less demanding evolutionary rationalizations or explanations of equilibrium behavior have become popular; see Sandholm (2010). This is the topic of Sections 3.2 and 3.3. Section 3.2 rehearses how NE can be reached by an evolutionary dynamics. That is, it does so in its reinterpretion as a deliberational dynamics, as suggested by Skyrms (1990) (Section 3.2). This interpretation is the more suitable one for us. In Section 3.3, then, we show how this account can be generalized to DE, and we shall propose two natural classes of dynamics that may converge to DE (that need not be NE), if they converge at all. The important upshot is: Evolutionary rationalizations apply to DE too.

Section 4, finally, is devoted to an extensive comparative discussion. Section 4.1 thoroughly compares CE and DE in order to reject the idea that CE are a sufficient treatment of the phenomenon of dependent beliefs. In Section 4.2, we argue that social norms and conventions are better captured by DE rather than by CE, as proposed by Vanderschraaf (1995) and Gintis (2009). Section 4.3 pursues the same aim vis à vis the suggestion of Binmore (2010) to explicate social norms via NE in in(de)finite repetitions of games. Quite a different approach to explaining cooperative behavior was put forward by Roemer (2019) with his idea of so-called Kantian equilibria. They have a moral touch. We shall see in Section 4.4 that Kantian equilibria are a special case of DE. Self-similarity need not be morally enforced, though, as conceived by Roemer. It seems to be quite a natural psychological tendency. Then, however, it is treated as the fallacy of magical thinking in the economic literature. This is discussed in Section 4.5, again with the aim of making clear that DE do not fall prey to this fallacy.

The motivation of program equilibria (Tennenholtz (2004)) and translucent equilibria (Halpern and Pass (2018)) resembles that of our entangled belief systems. They also implement the idea that players' decision procedures are mutually partially or fully transparent, so that players' actions are believed to be dependent. This is addressed in Sections 4.6 and 4.7. We shall see, however, that the formal explications diverge largely. Finally, we discuss the most suggestive idea that entangled belief systems rest on something like evidential reasoning. Al-Nowaihi and Dhami (2015) take up this idea and subsume it under the heuristics and biases program in cognitive psychology. As we shall see in Section 4.8, their notion of a consistent evidential equilibrium comes closest to our DE, but they are still not the same, and their interpretation is quite different. DE are not intended as an expression of bounded rationality. Section 4.9 finally explains that DE may find sympathy among evidential decision theorists, mainly found in philosophy. Just as NE build on causal decision theory, DE build on evidential decision theory. However, we emphasize that DE are not committed to this subsumption. As already indicated above, there is a causal story underlying entangled belief systems.

Section 5 provides a very brief conclusion.

This paper builds on Spohn (2003), where the idea of DE was first proposed. However, that paper had a different interest. It was largely concerned with causation- and decision-theoretic foundations, while the game-theoretic consequences were little worked out. It did not go beyond two-person games, a few examples, and a somewhat weaker version of Theorem 2 below. We still think that these foundations are important. However, in this paper we want to remedy this neglect and focus on the game-theoretic consequences. All in all, this paper is to demonstrate the significance, plausibility, and tractability of the novel notion of DE in various game-theoretic fields, to an extent showing that it deserves further attention and development.

2. Dependency Equilibria: Definitions, Examples, and Some Theorems

2.1 *Nash, Correlated, and Dependency Equilibria: The Completely Mixed Case*

Let us try to make precise the vague suggestions made so far, dealing only with games in normal form. We leave it open how to transfer our ideas to games in extensive form. Neither do we strive for maximal mathematical generalization. We are happy to deal with finitely many players pondering about finitely many pure strategies. So, let $I := \{1, \dots, m\}$ be the set of players, and S_i $(i = 1, \dots, m)$ be the set of pure strategies of player i. S_i is finite and has at least two members. Let $S := \times_{i \in I} S_i$ be the set of strategy profiles, with a typical element $s = (s_i, s_{-i})$, where $s_i \in S_i$ and $s_{-i} \in S_{-i} := \times_{j \neq i} S_j$. We speak of strategies here, as is common in game theory. However, it may be advisable to simply conceive of them as actions. The task of defining dependency equilibria for games in extensive form in which strategies can be detailed is beyond the present paper.

Each player i has a utility function u_i defined on the strategy profiles; i.e., $u_i : S \to \mathbb{R}$. Such a set-up $\gamma := \big(I, (S_i)_{i \in I}, (u_i)_{i \in I}\big)$ is called a game in normal form[4]; Γ denotes the set of all such games. Finally, let $\Delta(S)$ be the set of all probability distributions over S, $\Delta^+(S)$ the set of all completely mixed distributions $p \in \Delta(S)$ for which $p(s_i) = p(\{s_i\} \times S_{-i}) > 0$ for all $i \in I$ and $s_i \in S_i$,[5] and $\Delta_\perp(S)$ be the set of probability distributions $b \in \Delta(S)$ for which there are $b_i \in \Delta(S_i)$ for all $i \in I$ such that $b = \otimes_{i \in I} b_i$ (i.e., in b each b_i is probabilistically independent from all the other b_j; the letter b is to signal this).[6]

For comparison, let us rehearse the traditional notions:

Definition 1. Let $\gamma \in \Gamma$. Then $b \in \Delta_\perp(S)$ is a NE in γ iff for all $i \in I$, $s_i \in S_i$ with $b(s_i) > 0$, and $s'_i \in S_i$,

$$\sum_{s \in S} u_i(s) b(s) = \sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) b(s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} u_i(s'_i, s_{-i}) b(s_{-i}). \tag{1}$$

Here, we may interpret b_i as a mixed strategy of player i. In a NE player i can do no better in terms of expected utility than choosing the mixed strategy b_i, given the other players stick to their mixed strategies in the NE. Or we may interpret a NE b as an equilibrium of beliefs. Then b_i expresses the expectation the other players have about player i. And the NE b expresses the idea that the choices of the players are believed to be probabilistically independent and that, given her expectations about the other players, each player can

[4] Throughout this text γ always represents a tuple $\big(I, (S_i)_{i \in I}, (u_i)_{i \in I}\big)$.

[5] Note that this is slightly weaker than requiring that p is strictly positive or regular. For a completely mixed $p \in \Delta^+(S)$ we might still have $p(s) = 0$ for some $s \in S$.

[6] If $b \in \Delta_\perp(S)$ is completely mixed, then, of course, b is strictly positive.

do no better in terms of expected utility than taking one of the pure strategies ascribed positive probability by the other players. For us, this will be the preferred interpretation.

The independence assumption is given up in the idea of correlated equilibria first ventured by Aumann (1974). Any distribution $p \in \Delta(S)$ can be a CE. To be precise, two definitions are in use, as observed by Aumann (1987) (p. 12). Right now, we present only the definition of what is called a CE distribution (by Aumann) or a canonical CE, because it is directly comparable with the definition of NE and that of DE to be given. Aumann's original definition is a different one. We shall return to this point in Section 3.1.

DEFINITION 2. Let $\gamma \in \Gamma$. Then $p \in \Delta(S)$ is a (canonical) CE in γ iff for all $i \in I$, $s_i \in S_i$ with $p(s_i) > 0$, and $s_i' \in S_i$,

$$\sum_{s \in S} u_i(s)p(s) = \sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})p(s_{-i}|s_i) \geq \sum_{s_{-i} \in S_{-i}} u_i(s_i', s_{-i})p(s_{-i}|s_i). \tag{2}$$

Here, p can be interpreted as a joint mixed strategy profile. Myerson (1991) (p. 253) suggested that a mediator plays out this mixture and informs each player only about her assignment. Should a player accept her assignment? After the assignments, p results in a certain expected utility for each player i. If, for all players i, this is not exceeded by the expected utility of any of i's pure strategies with respect to the probabilities $p(s_{-i}|s_i)$ conditional on the mediator's assignment of s_i, then, and only then, p is a correlated equilibrium in which no player has a reason to abolish the dependence. This may be interpreted as an equilibrium of beliefs, too. The beliefs of each player i about what the others will do all derive from one joint distribution *p* and the mediator's assignment of s_i. Then, in a CE p, abolishing the dependence and choosing independently of the others cannot be better in terms of expected utility than accepting the dependence and playing out p by following the proposal.

We want to suggest that this is not the only standard of comparison and perhaps not the most suitable one. As the second and third term of (2) make clear, in a CE we compare the conditional expected utilities of the s_i mixed by $p(s_i)$ with the unconditional expected utilities of the s_i' after the mediator's assignment of s_i. However, once we allow probabilistic dependence between choices and 'states of the world', i.e., the other players' choices, optimization is always about conditional expected utilities. This is how we roughly explained the idea of dependency equilibria in the introduction.

DEFINITION 3. Let $\gamma \in \Gamma$. Then $p \in \Delta^+(S)$ is a completely mixed (canonical) DE in γ iff for all $i \in I$, $s_i \in S_i$ with $p(s_i) > 0$, and $s_i' \in S_i$,

$$\sum_{s \in S} u_i(s)p(s) = \sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})p(s_{-i}|s_i) \geq \sum_{s_{-i} \in S_{-i}} u_i(s_i', s_{-i})p(s_{-i}|s_i').^{7} \tag{3}$$

Obviously, this definition works only in the completely mixed case where $p \in \Delta^+(S)$. Otherwise we might have $p(s_i') = 0$ so that the last term of (3) is undefined. Getting rid

[7] This definition was first proposed in Spohn (2003) (p. 200), though only for two-person games.

of this restriction will require some work beyond the standard probabilistic framework. We leave this to Section 2.2. Still, let's stick for a moment to the restricted case, because it eases the understanding of the basic idea of DE and its comparison with NE and CE.

Let us observe right away:

PROPOSITION 1. *For all $\gamma \in \Gamma$, each NE in γ is a CE, and each completely mixed NE is a completely mixed DE in γ, but not vice versa. And there may be completely mixed CE in γ which are not completely mixed DE in γ, and vice versa.*

As to the first claim about CE, Aumann (1974) (p. 78) has proved that the convex closure of all NE in γ is a subset of the set of all CE in γ and often a proper one. The first claim about DE is evident. The second claim of Proposition 1 will be verified by examples in Section 2.3. There we will also see that the restriction of proposition 1 to the completely mixed case is irrelevant. Thus, CE and DE are two independent generalizations of NE.

The difference between Definitions 2 and 3 immediately catches the eye. Definition 2 represents optimization by breaking the dependency after receiving the mediator's assignment, while Definition 3 represents optimization within the dependency. No player can do better in terms of conditional expected utility by changing to a strategy not receiving positive probability in the equilibrium. Again, one might interpret this as a choice of a joint mixed strategy combination, e.g., administered by a mediator, to which all players agree. However, it is perhaps more natural to interpret a DE as an equilibrium of beliefs. Each player has possibly varying probabilities for the choices of the other players conditional on her own choices, as dictated by the DE p, and relative to them she maximizes her conditional expected utilities by choosing one of her options having positive probability in p. The relation between CE and DE will be discussed in more detail in Section 4.1.

In Section 1 we have indicated a causal story as to how agents' beliefs may take such a shape. We will return to it and related stories in Section 3. Section 3.3 will add an evolutionary story. Whatever the reasons for the players to be in a state of entangled belief systems, DE provide the appertaining equilibrium concept. Let us study them more carefully.

2.2 *Dependency Equilibria: The General Case and Some Observations*

Due to its restriction to completely mixed distributions, Definition 3 of DE was not yet satisfactory. Of course, we must not confine DE to such distributions. Let's get this straight in this section. The obvious idea is to approximate the general case by completely mixed distributions. This leads to

DEFINITION 4. Let $\gamma \in \Gamma$. $p \in \Delta(S)$ is a limit dependency equilibrium (limDE) iff there is a sequence $(p_r)_{r\in\mathbb{N}}$ with $p_r \in \Delta^+(S)$ for all $r \in \mathbb{N}$ such that $\lim_{r\to\infty} p_r = p$ and, for all $i \in I$ and $s'_i \in S_i$,

$$\lim_{r\to\infty} \sum_{s\in S} u_i(s)p_r(s_{-i}|s_i)p_r(s_i) \geq \lim_{r\to\infty} \sum_{s_{-i}\in S_{-i}} u_i(s'_i, s_{-i})p_r(s_{-i}|s'_i).\text{[8]} \tag{4}$$

[8]This definition is also given in Spohn (2003) (p. 201) for two-person games.

This allows for the possibility that the left terms are smaller than the right terms for all $r \in \mathbb{N}$ and reach equality only in the limit. Perhaps one should exclude this possibility. However, we don't pursue here this theoretical option.

We should add right away:

DEFINITION 5. $p \in \Delta(S)$ is a proper DE in γ iff p is a limDE, but not a NE in γ. p is a pure DE in γ iff p is a limDE in γ and there is an $s \in S$ with $p(s) = 1$. Otherwise, p is a mixed DE in γ. And p is a completely mixed DE in γ iff moreover $p \in \Delta^+(S)$.

There is another, perhaps more straightforward way to correct Definition 3. We may assume a probabilistic structure in which probabilities conditional on null propositions are defined. This is specified in

DEFINITION 6. P is a Popper measure for S iff $P(A|B)$ is defined for all $A \subseteq S$ and all B with $P(B|C) > 0$ for some $C \subseteq S$, such that $P(.|B)$ is a probability measure for S and $P(A \cap B|C) = P(A|B \cap C)P(B|C)$, whenever $P(A|B \cap C)$ is defined. Let $\Delta_{Popp}(S)$ be the set of all Popper measures for S. We will write $P(A)$ for $P(A|S)$.

We use "P" in order to denote this kind of measure. The label honors Popper (1938) where such a structure was first conceived. There are variations concerning the set of propositions for which conditional probabilities are defined (see, e.g., Spohn (1986)).

The notion of lexicographic probabilities is perhaps a more familiar way of dealing with null conditions:

DEFINITION 7. A sequence $\lambda := (\lambda_0, \ldots, \lambda_q)$ is a lexicographic probability for S iff each λ_k ($k \leq q$) is a probability measure for S such that $\lambda_{k+1}\big(\{s \in S : \lambda_l(s) = 0 \text{ for all } l \leq k\}\big) = 1$. Let $\Delta_{lex}(S)$ be the set of all such sequences.

However, the two notions are equivalent:

LEMMA 1. *Let $\lambda \in \Delta_{lex}(S)$. Define P by $P(A|B) = \lambda_k(A|B)$ where $k = \min_{l \leq q}\big(\lambda_l(B) > 0\big)$. Then $P \in \Delta_{Popp}(S)$. Conversely, let $P \in \Delta_{Popp}(S)$. Define $\lambda_0 := P(.|S)$, and if $T_k := \{s \in S : \lambda_k(s) > 0\}$, define $\lambda_{k+1} := P(.|S \setminus \bigcup_{l \leq k} T_l)$. Finally, let q be the largest $k+1$ for which $P(.|S \setminus \bigcup_{l \leq k} T_l)$ is still defined. Then $\lambda = (\lambda_0, \ldots, \lambda_q) \in \Delta_{lex}(S)$. Thus, if $T = \bigcup_{l \leq q} T_l$, P(A | B) is defined iff $B \cap T \neq \emptyset$ and undefined iff $B \subseteq S \setminus T$.*

PROOF. See van Fraassen (1976), and Spohn (1986) for the σ-additive case. □

In our context, we must specifically consider such Popper measures for which all probabilities conditional on the players' strategies are defined. Hence, let $\Delta^*_{Popp}(S)$ be the set of Popper measures P for S for which $P(.|s_i)$ is defined for all $s_i \in S_i$, $i \in I$. Then we can straightforwardly define dependency equilibria in terms of these notions without recourse to limit constructions:

DEFINITION 8. Let $\gamma \in \Gamma$. Then P is a lexicographic dependency equilibrium (lexDE) iff $P \in \Delta^*_{Popp}(S)$ and for all $i \in I$ and $s'_i \in S_i$,

$$\sum_{s \in S} u_i(s) P(s_{-i}|s_i) P(s_i) \geq \sum_{s_{-i} \in S_{-i}} u_i(s'_i, s_{-i}) P(s_{-i}|s'_i). \tag{5}$$

Note that equation (5) is identical with (3) except that the standard *p* is replaced by the Popper measure P. Thus, Definition 8 is the straightforward generalization of Definition 3 we looked for.

Now we have two generalizations of Definition 3, a lexicographic one and one in terms of limits. What is their relation? They cannot be strictly equivalent because lexDE (Popper measures) are more fine-grained than limDE (standard distributions). Still, they are essentially equivalent. We defer showing this to Appendix A.

How to compute DE? In Section 2.3, we shall consider some examples of 2×2 two-person games. There, DE are easy to calculate. If there are only (completely) mixed DE, then the weak inequalities in Definition 3 are in fact equations. Hence, we have three equations for the four unknowns in the probability matrix. Therefore, we usually have a one-dimensional family of DE in 2×2 two-person games. We will find, though, that when computing DE in 2×2 two-person games one already gets entangled in quadratic equations. And it is easily seen that the order of the polynomial equations increases with the number of strategies available to the two players. This is in sharp contrast to the computation of NE in two-person games with any finite number of strategies, which requires only linear programming for solving linear equations.

However, this observation is deceptive. Datta (2003) proved that finding Nash equilibria in three-person games is equivalent to finding the solutions of any polynomial equation system whatsoever, where the number of strategies of the three players required for this equivalence is effectively specifiable. He also showed the same equivalence for *n*-person games in which each player has just two strategies and effectively specified the number *n* of players. So, as soon as we leave the domain of two-person games the solution theory of Nash equilibria becomes arbitrarily complex as well.

Generally, we can say that the solution space of DE is high-dimensional. If player $i \in I = \{1, \ldots, m\}$ has $n_i \geq 2$ strategies in S_i, then there are $N := \prod_{i \in I} n_i$ strategy profiles in S, and an equal number of probability values has to be specified in a limDE. However, in a completely mixed limDE we have only $\sum_{i \in I} n_i - m + 1$ equations for determining these values. In case of not completely mixed DE, inequalities replace some of the equations. For some first results on the solution theory for two-person games see Portakal and Sturmfels (2022). If the notion of a DE is indeed an interesting one, then the interesting parts of this solution space should be identified and more deeply investigated.

For the moment, however, we can only offer, as it were, lower and upper bounds on this solution space. The greatest lower bound is specified by the security levels of the players, where the *security level* $\underline{u}_i$ of player i is defined as $\underline{u}_i = \max_{s_i \in S_i} \Big(\min_{s_{-i} \in S_{-i}} \left(u_i \left(s_i, s_{-i} \right) \right) \Big)$. More precisely we have:

PROPOSITION 2. *Let* $\gamma \in \Gamma$ *and* $p \in \Delta(S)$. *If* $\sum_{s \in S} u_i(s) p(s) < \underline{u}_i$ *for some player* i, *then* p *is not a limDE.*

PROOF. Let P be a Popper measure with $p = P(.|S)$. If P were a lexDE, then we should have $\underline{u}_i > \sum_{s \in S} u_i(s)P(s) \geq \sum_{s_{-i} \in S_{-i}} u_i(s'_i, s_{-i})P(s_{-i}|s'_i)$ for all $s'_i \in S_i$. But this is impossible; we must have $\sum_{s_{-i} \in S_{-i}} u_i(s'_i, s_{-i})P(s_{-i}|s'_i) \geq \underline{u}_i$ at least for that s'_i for which $u_i(s'_i, s_{-i}) = \underline{u}_i$. And if P can't be a lexDE, p can't be a limDE. $\square$

Hence, no player can fall below her security level in a DE. However, players can fall so low. The 'bad' DE in the Hi-Lo game discussed in Section 2.3 will demonstrate that there are DE in which all players realize no more than their security levels, even though they could do much better.[9] Note also that Proposition 2 only says that there cannot be a DE in which the players' expectations are below their security level. This still allows that strategy profiles falling below their security levels are part of a DE. We will find this situation exemplified in the ultimatum game discussed in Section 2.3.

There is also a kind of lowest upper bound for DE which is more interesting. It is given by the weakly Pareto-optimal strategy profiles in S. More precisely, we have:

THEOREM 1. *Let p be a limDE and let $s^* \in S$ be at least as good as p in the sense that for all $i \in I$, $u_i(s^*) \geq \sum_{s \in S} u_i(s)p(s)$. Then p^* defined by $p^*(s^*) = 1$ is a pure limDE.*

PROOF. The proof is best given in terms of Popper measures and lexDE. First, assume that $p(s^*) > 0$. Then s^* is in the support of the (mixed) DE p, so that $u_i(s^*) = \sum_{s \in S} u_i(s)p(s)$ for all $i \in I$. Thus, p^* is equivalent to the DE p and itself a (pure) DE.

So, assume $p(s^*) = 0$ instead. We want to show that there is a $P^* \in \Delta^*_{Popp}(S)$ such that $p^* = P^*(.|S)$ and P^* is a lexDE. Since p is a limDE, there must be a lexicographic probability $\lambda = (\lambda_0, \ldots, \lambda_q)$ forming a corresponding lexDE. Let P be the corresponding Popper measure. Now define a new lexicographic probability $\lambda^* = (p^*, \lambda_0, \ldots, \lambda^*_k, \ldots, \lambda_q)$, where k is the smallest index for which $\lambda_k(s^*) > 0$ and $\lambda^*_k = \lambda_k(.|S \backslash \{s^*\})$. And define P^* to be the Popper measure corresponding to λ^*. Then $P^* \in \Delta^*_{Popp}(S)$, since the required conditional probabilities are defined in P^*. And P^* is indeed a lexDE, as we show now.

First, we have for all $i \in I$ and $s'_i \neq s^*_i$,
$\sum_{s \in S} u_i(s)P^*(s_{-i}|s_i)P^*(s_i) = u_i(s^*)$ (because $P^*(s^*) = 1$)
$\geq \sum_{s \in S} u_i(s)p(s)$ (by assumption)

[9]One might think that being a strategy profile that is at least as good as the players' security levels is sufficient for being a pure DE. However, this conjecture is wrong. Consider the following two-person game:

2,2	3,3	-1,-1
3,3	0,5	5,0
-1,-1	5,0	0,5

Here, (2, 2) is not an NE. There is only one mixed NE, namely, the equal distribution over the lower right four fields, with an expected utility of 2.5 for each player. The players' maximin strategies are their second and third strategies, their security levels being 0. So, if this conjecture would hold, (2, 2) should be a DE. But it isn't. However we distribute the probabilities over the lower right four fields, the conditional expected utility of at least one player is $\geq$ 2.5. Giving weight to (3, 3) increases conditional expected utility, while (-1, -1) is below the security levels and receives weight 0. So, this is a counter-example to the conjecture.

$\geq \sum_{s_{-i}\in S_{-i}} u_i(s_i', s_{-i})P(s_{-i}|s_i')$ (because P is lexDE corresponding to the limDE p)
$= \sum_{s_{-i}\in S_{-i}} u_i(s_i', s_{-i})P^*(s_{-i}|s_i')$,
because either $P^*(s_{-i}|s_i') = \lambda_l(s_{-i}'|s_i)$ for $l \neq k$ or $P^*(s_{-i}|s_i') = \lambda_k^*(s_{-i}|s_i') = \lambda_k(s_{-i}|s_i'\backslash s^*) = \lambda_k(s_{-i}|s_i')$, since $s_i' \neq s_i^*$ (here s_i', etc. are considered as subsets of the space S).

It remains to be shown that the same holds for $s_i' = s_i^*$. As before we have: $u_i(s^*) \geq \sum_{s\in S} u_i(s)p(s) \geq \sum_{s_{-i}\in S_{-i}} u_i(s_i^*, s_{-i})P(s_{-i}|s_i^*)$. And $P(s_{-i}|s_i^*) = \lambda_l(s_{-i}|s_i^*)$, where l is the smallest index for which $\lambda_l(s_i^*) > 0$. Clearly, $l \leq k$.

If $l \neq k$, then $\lambda_l(s_{-i}|s_i^*) = \lambda_l^*(s_{-i}|s_i^*) = P^*(s_{-i}|s_i^*)$, and hence
$u_i(s^*) \geq \sum_{s_{-i}\in S_{-i}} u_i(s_i^*, s_{-i})P^*(s_{-i}|s_i^*)$.

If $l = k$, we have
$u_i(s^*) \geq \sum_{s_{-i}\in S_{-i}} u_i(s_i^*, s_{-i})p(s_{-i}|s_i^*) = \sum_{s_{-i}\in S_{-i}} u_i(s_i^*, s_{-i})\lambda_k(s_{-i}|s_i^*)$
$= u_i(s^*)\lambda_k(s_{-i}^*|s_i^* + \sum_{s_{-i}\neq s_{-i}^*} u_i(s_i^*, s_{-i})\lambda_k(s_{-i}|s_i^*)$. Subtracting the first term on both sides and then dividing by $1 - \lambda_k(s_{-i}^*|s_i^*) = \lambda_k(s_i^*\backslash s_{-i}^*|s_i^*) = \lambda_k(s_i^*\backslash s^*|s_i^*)$, we get:
$u_i(s^*) \geq \sum_{s_{-i}\neq s_{-i}^*} u_i(s_i^*, s_{-i})\lambda_k(s_{-i}|s_i^*)/\lambda_k(s_i^*\backslash s^*|s_i^*)$
$= \sum_{s_{-i}\neq s_{-i}^*} u_i(s_i^*, s_{-i})\lambda_k(s_{-i}|s_i^*\backslash s^*) = \sum_{s_{-i}\neq s_{-i}^*} u_i(s_i^*, s_{-i})\lambda_k^*(s_{-i}|s_i^*)$
$= \sum_{s_{-i}\in S_{-i}} u_i(s_i^*, s_{-i})\lambda_k^*(s_{-i}|s_i^*)$ (because $\lambda_k^*(s_{-i}^*|s_i^*) = 0$)
$= \sum_{s_{-i}\in S_{-i}} u_i(s_i^*, s_{-i})P^*(s_{-i}|s_i^*)$. □

Theorem 1 entails that each weakly Pareto-optimal strategy profile which is at least as good as some DE constitutes a DE in turn.[10] Proposition 2 specifies a necessary condition for a strategy profile to constitute a pure DE, and Theorem 1 a sufficient condition. We don't know of a condition that is both necessary and sufficient for pure DE.

We take Theorem 1 to be a significant observation. Pareto-optimality is commonly taken as the basic (and quite weak) criterion of social optimality or collective rationality (if we only knew what this precisely is). The alarming point then is that individual rationality may apparently conflict with social optimality, as paradigmatically demonstrated by Prisoners' Dilemma. According to Theorem 1, however, this opposition does not obtain. If DE are individually rational for the players by maximizing their conditional expected utility, then social or Pareto-optimality is always in reach of individual rationality. This will clearly be exemplified by our discussion of Prisoners' Dilemma below, where the cooperative solution turns out to be a pure DE even in the one-shot game. Pareto-optimality is thus at least attainable, though there may be many non-Pareto-optimal DE too. How to choose among Pareto-optimal strategy profiles remains as open as the general problem of equilibrium selection.

2.3 *Some Examples*

Let us illustrate the novel equilibrium concept with some examples in order to get a feeling for how fundamentally DE differ from NE and also CE. We choose some 2×2 two-person games exemplifying typical social situations. Spohn (2003) illustrates DE

[10] In Spohn (2003) (pp. 208f., Observation 5) only a special case of Theorem 1 was proven, namely only for two-person games and only for strategy combinations that are at least as good as some NE. The proof proceeded only in terms of limDE.

with Matching Pennies, Bach or Stravinsky, Hawk and Dove (= Chicken), and Prisoners' Dilemma. We choose here Hi-Lo, the Stag Hunt, and the Ultimatum Game, which are perhaps more illustrative, while also repeating Prisoners' Dilemma. Throughout, we shall denote the strategies of player 1, the row chooser, by a_1 and a_2 and the strategies of player 2, the column chooser, by b_1 and b_2.

	b_1	b_2
a_1	$2, 2^{*,+}$	$0, 0^{+}$
a_2	$0, 0^{+}$	$1, 1^{*,+}$

FIGURE 1. HL. Pure NE are marked with $*$ and pure DE with $+$.

p	b_1	b_2
a_1	1/9	2/9
a_2	2/9	4/9

FIGURE 2. HL. Mixed NE

p	b_1	b_2
a_1	0	$x \in [0, 1]$
a_2	1-x	0

FIGURE 3. HL. 1st family of DE

p	b_1	b_2
a_1	$x \in [0, 1/9]$	$\frac{(1-x-y)}{2}$
a_2	$\frac{(1-x-y)}{2}$	$y = \frac{(1-x \pm \sqrt{(1-x)(1-9x)})}{2}$

FIGURE 4. HL. 2nd family of DE

The *Hi-Lo* game (*HL*, Figure 1): This is the paradigmatic coordination game (with unequal gains). It has two pure and one mixed NE (Figures 1 and 2). Its CE are the two pure NE, all mixtures of them, and many more, indeed all distributions where both probabilities on the counterdiagonal are not greater than twice the probability of (a_1, b_1) and half of the probability of (a_2, b_2).

Its DE, however, have an astonishingly complex structure: First, the two pure NE are also pure DE (Figure 1). Then there is the family of proper DE depicted in Figure 3. Here, both players are caught in an unhappy dependence. Whatever the one player does, the other is guaranteed to do the opposite, and both have a conditional expected utility of 0. For $x = 1$ or 0, these are indeed further pure DE (Figure 1). In analogy to NE, we might define strict DE and then observe that, whatever x, none of these DE are strict. Thus, these DE would be excluded, if we would postulate that, rationally, players should go for a strict DE, if there is one. However, we shall be silent here on rationality postulates beyond maximizing conditional expected utility. Note also that the example shows the missing part of Proposition 1, i.e., the independence of CE and DE. The DE of Figure 3 are not CE, and most of the CE offside the diagonal are not DE.

Finally, the mixed DE unfold in another family of DE, or in fact two families (Figure 4). For each admissible x, y takes two admissible values according to the $\pm$. In all these cases, the conditional expected utilities for both strategies are the same. Again, we might ponder about an additional rationality postulate, namely that players should choose a weakly Pareto-optimal DE. In finite games, they are guaranteed to exist. If there are many, this postulate is of limited help. But in our case there is only one, and the problem of equilibrium selection would be solved by this postulate (as with NE). Self-similarity considerations (see Section 4.5) and team reasoning, as suggested by Sugden

(1993) would come to the same conclusion in this case. However, as mentioned, we do not treat the problem of equilibrium selection in this paper. It is a central, huge and much treated problem for NE. Here, we can only observe that it looks still more pressing for DE; but we have seen in Section 2.2 that it is also quite different.

Note also that mixtures of the pure NE of *HL* are CE, but not DE. This changes, however, when we consider *HL* with equal payoffs (this is the proper coordination game). It is easily verified that in this case mixtures of pure NE are also DE. So, whatever arguments and evidence speak for CE in these games (as capturing social norms – see Section 3.3), they also speak for DE.

The *Prisoners' Dilemma* game (*PD*, Figure 5): This vigorously debated game is about an alleged conflict between individual and collective rationality, or about an alleged impossibility of cooperation. It comes in many guises, also for n-person games, e.g., as the free-rider problem or the tragedy of the commons. It is perhaps the most famous example of game theory.

	b_1	b_2
a_1	$2, 2^+$	$0, 3$
a_2	$3, 0$	$1, 1^{*,+}$

FIGURE 5. PD. Pure NE are marked with $*$, pure DE with $+$.

p	b_1	b_2
a_1	$\frac{x(1+x)}{2}$	$\frac{x(1-x)}{2}$
a_2	$\frac{x(1-x)}{2}$	$\frac{(1-x)(2-x)}{2}$

FIGURE 6. PD. 1st family of DE

p	b_1	b_2
a_1	$\frac{(1-x)(1+x)}{8}$	$\frac{(1-x)(1-3x)}{8}$
a_2	$\frac{(1-x)(1+3x)}{8}$	$\frac{(1-x)(1+x)}{8}$

FIGURE 7. PD. 2nd family of DE

Not only is joint defection (a_2, b_2) the only NE and indeed the only CE; the force of *PD* is that for each player defection strictly dominates cooperation.

By contrast, we have again two families of DE, a symmetric family (Figure 6) and an asymmetric family (Figure 7). In the second family, a likely cooperator stands in an unhappy dependence with a likely defector. However, the bias cannot get extreme. In the extreme, the unhappy cooperator would receive payoff 0, but in a DE her expectation can be no less than her security level 1, as guaranteed by Proposition 2.

The first, symmetric family is still more interesting. For $x = 0$ we get the NE of two defectors. However, for $x = 1$ we get another pure DE of two cooperators. So, if we can rationalize DE, we can rationalize pure cooperation in *PD*, at least as a rational possibility. Again, if we add the rule to rationally choose a Pareto-optimal DE among the available DE, cooperation even becomes a rational necessity.

The invention of DE originated from this point. There are numerous attempts at avoiding the dilemma and rationalizing cooperation. Most of them work by changing the game structure (e.g., through allowing binding preplay communication or exit moves) or by changing the utility functions (e.g., through sanctioning defection), or by moving into an iterative or sequential context (e.g., in terms of sequential equilibria or folk theorems). So, do we need another rationalization of cooperation? There is no point now in entering a comparative discussion. Certainly, these attempts are all instructive. Spohn (2003) (pp. 196f.), however, has expressed dissatisfaction with the existing ideas and proposed to grab the problem at its root, the unaltered single-shot *PD*, and not merely to dissolve it by altering the game. DE provide a way to do so. Insofar, DE offer a novel treatment of *PD*.

There are various other attempts. E.g., constrained maximization, introduced and defended by Gauthier (1986) results in joint cooperation in *PD*; but his account is plagued by the issue of when and why to be a constrained rather than a straightforward maximizer. Another idea establishing joint cooperation in the unaltered single-shot *PD* is team reasoning; see Sugden (1993, 2011) and Karpus and Radzvilas (2018). However, they also face the issue when to rationally recommend team instead of best-response reasoning. In Section 3 we will find still further attempts. We could go on here for long. In any case, our rationalization of joint cooperation in PD is only as good as the rationalization of DE themselves. We will attend to this issue more carefully in Section 4.

Let us finally look at a game which presents quite an artificial situation, but which has attracted the attention of behavioral economists because it allegedly shows another dramatic inadequacy of orthodox game theory.

The *Ultimatum* game (*UG*, Figure 8): We present a minimal version of it, in order to keep to our format of 2 x 2 two-person games. The proponent can divide \$4 either 3 to 1 in his favor or fairly 2 to 2. The respondent can accept the proposal or reject it; in the latter case nobody gets anything:

	b_1	b_2
a_1	$3, 1^{*,+}$	$0, 0$
a_2	$2, 2^{+}$	$0, 0$

FIGURE 8. UG. Pure NE are marked with $*$, pure DE with $+$.

Of course, the only NE, and the only CE, is (a_1,b_1). The proponent takes \$3 and the respondent accepts. For both, this is even the strictly dominant choice.[11] But the NE is rarely observed in experiments. More often, the proponent offers a roughly fair division (in a fuller version of the game), without accepting the charge of irrationality. So, something seems wrong with NE, either as a recommendation or as a prediction. The common conclusion is that orthodox game theory needs to be amended by something like fairness considerations.

This conclusion is not imperative, though. For, what are the pureDE of *UG*? First, the above NE. But also fair division: when the probability of rejection is high enough given

[11] *UG* is usually presented in extensive form. Then the respondent sees the proponent's division. In the normal form this feature of *UG* vanishes. As long as we do not extend the notion of DE to games in extensive form, this slight misrepresentation is unavoidable.

the unfair proposal and 0 given the fair proposal. So, given the players are in a suitable dependence, fairness results from a DE. As in *PD*, we can establish the feasibility of a satisfactory solution without adding any considerations extraneous to the game. This deems us a virtue of our account. Note, however, that the emphasis on fairness is generated only by our minimal version with just one unfair and one fair outcome (besides disaster). In a fuller version, there are many possible divisions of the amount in play, and each division would constitute a DE. Still, (roughly) fair division is among them.

This discussion of some basic examples as well as the beginning of a theory in Section 2.2 show that considerable changes are forthcoming when we build game theory on the notion of DE instead of NE. The changes certainly call for further elaboration.

3. Rationalizations

We have emphasized that DE firmly stand in the tradition of best-response reasoning. Therefore, they are able not only to explain, but indeed to rationalize human behavior. In this section, we want to elaborate on the rational justification of DE in the ways well-established for NE. That is, in Section 3.1 we show how DE can be derived from common knowledge assumptions, and, after necessary preparations in Section 3.2, we explain in Section 3.3 how DE can emerge from evolutionary processes. So, the message is that these well-known paths are not restricted to NE, but open to DE as well.

3.1 *Dependency Equilibria within Epistemic Game Theory*

As said, we maintain that DE can be defended from the perspective of individual rationality and not only from that of social desirability. So far we have tacitly shared the common assumption that NE are sufficiently well justified from the individual perspective, and we have argued that DE are the most suggestive generalization of NE within the context of entangled belief systems. And then we implied that the individual rationality of NE somehow transfers to their generalization to DE. In this section we want to vindicate this implication.

The most common justification of NE is in terms of public advice. When one gives public advice to players indicating which (mixed) strategies to choose, one can only recommend a NE; otherwise, the advice would be self-defeating. This argument about offering a 'public roulette' was also used by Aumann (1974) (p. 84) to justify CE. If the offer of playing out a joint probability distribution is not to be self-defeating, it must not give any player a reason to deviate — in terms of posterior *unconditional* expected utility. This suggests, however, that the very same argument applies to DE as well. If you propose a joint probability distribution, it must be a DE. Otherwise, the proposal is self-defeating because it gives positive probability to a strategy of at least one player that has less than maximal *conditional* expected utility.

From epistemic game theory we have learned that the story about public advice is not required for justifying equilibrium choice. The essential point about the publicity of advice is that it generates *common knowledge*, as first explicated by Lewis (1969) (sect. II.1). Common knowledge of players' rationality and of their probabilities (as well as of the game structure and utilities) suffices for deriving a NE.[12] Note that a NE is thereby presented as an epistemic equilibrium among the mutual beliefs of the players about each other. In this way, game-theoretic rationality is reduced to decision-theoretic rationality plus common knowledge assumptions. This was indeed the goal of epistemic game theory. It does not matter how this common knowledge comes about. Public advice is one good way of establishing it. But there may be other ways as well. It may sometimes be established by reasoning; there may be a joint learning history in the background; etc. We will tell a story about how such common knowledge may at least be approximated in Sections 3.2 and 3.3 on deliberational dynamics.

Let us transfer this well-known story to DE in formal detail. The first thing to fix is the format of the players' first-order beliefs. We take mutual knowledge of the game structure and utility functions for granted. The first-order beliefs of player i about strategy profiles split into two parts. There is, first, i's indecision state $p_i \in \Delta(S_i)$ concerning her own strategies. And there is, secondly, i's assessment $p_{-i}(s_{-i}|s_i)$ of the other players' strategies s_{-i} given her own choice s_i. These two parts combine into a full joint distribution $p \in \Delta(S)$. It will turn to be conceptually desirable to distinguish the two parts, when we extend our study of deliberational dynamics to DE in Section 3.3. However, presently it is less cumbersome to assume the first-order beliefs of player i just to consist in the full distribution $p \in \Delta(S)$ without distinguishing its parts.

If we proceed with this format, we would have to somehow relate these first-order beliefs to limDE. It is not clear, though, how to do so, since we have no epistemic interpretation of the only mathematically motivated sequence of positive distributions converging to a limDE. Hence, our discussion better proceeds in terms of lexDE. But then we should rather conceive of the first-order beliefs of the players as a full Popper measure $P \in \Delta^*_{Popp}(S)$, i.e., as something for which the notion of a DE is directly explained.

Next, we must assume a vector $Q := (Q_1, \dots, Q_m)$ of second-order beliefs of the players $i \in I$ about the other players. If first-order beliefs are Popper measures, second-order beliefs Q_i should presumably be so as well. We shall see, though, that only the first member of the corresponding lexicographic probability is relevant. Q_i contains a distribution over the first-order beliefs of $j (j \neq i)$, i.e., over the set of Popper measures P_j that j might have. However, we assume that $Q = (Q_1, \dots, Q_m)$ has mutual knowledge of the Popper measure $P \in \Delta^*_{Popp}(S)$ in the sense that for each $i \in I$, $P_i = P$ and Q_i(j has first-order beliefs $P) = 1$ for all $j \neq i$. This avoids the potential complexity of second-order beliefs. What does the Popper measure Q_i say, given the null condition that j does not have first-order beliefs P? Here we might assume anything, even that it does not say anything at all. It will turn out to be irrelevant.

What does mutual knowledge of rationality mean in the present setting? Well, $s_i \in S_i$ is a best response to first-order beliefs $P_i \in \Delta^*_{Popp}(S)$ iff s_i maximizes i's conditional expected utility with respect to P_i, i.e., iff $\sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) P_i(s_{-i}|s_i)$ is maximal within S_i.[13] And $Q = (Q_1, \dots, Q_m)$ has mutual knowledge of rationality iff for each $i \in I$, Q_i(j

[12]If one merely assumes common knowledge of rationality, one is automatically led to the notion of rationalizability; see Spohn (1982) (pp. 251f.), Bernheim (1984), and Pearce (1984).

[13]It may sound odd to call s_i a best response to the full P_i. It seems more appropriate to say that s_i is a best response to the collection of all conditional probabilities $P_i(s_{-i}|s'_i)$. However, we have concluded not to

takes some best response to her first-order beliefs $P_j) = 1$ for all $j \neq i$. Again, the question arises: what do the Q_i believe given another player does not choose a best response? Again the answer turns out to be irrelevant. We may assume anything.

This simplifies. If $Q = (Q_1, \ldots, Q_m)$ has mutual knowledge of the Popper measure P, then each player i knows the first-order beliefs P of each other player j. And mutual knowledge of rationality reduces to the fact that for each $i \in I$, Q_i agrees with P in the sense that $P(s_j) = 0$ if $Q_i(s_j$ is not a best response to $P) = 1$ for all $j \neq i$. Then we have:

THEOREM 2. *If, in the game γ, $Q = (Q_1, \ldots, Q_m)$ has mutual knowledge of rationality and mutual knowledge of the distribution $P \in \Delta^*_{Popp}(S)$, then P is a lexDE.*

PROOF. LexDE are defined by condition (5), which holds for all $i \in I$ and $s_i' \in S_i$, because only those strategies s_i of i which maximize conditional expected utility have weight $P(s_i) > 0$ in the mixture on the LHS of (5). □

This theorem is as trivial as in the case of NE. It does not present an exciting story. But it shows that whatever rationale we have for NE, we can directly duplicate it for DE. We could, it seems, develop epistemic game theory also in terms of DE. Of course, one might object how unrealistic the mutual knowledge assumptions are. On the other hand, recall how easily such assumptions are generated by public announcements. Be this as it may, more realistic stories need not be bound to be radically different. In the next section we shall work towards a bit more realism by putting DE into the context of evolutionary or rather of deliberational dynamics.

3.2 *Deliberational Dynamics and Nash Equilibria*

Another justification of NE with less stringent presuppositions is provided by evolutionary game theory. There it is shown how the evolution of a population consisting of boundedly rational agents using simple strategy revision protocols may converge to a NE in a one-shot non-cooperative game (for a comprehensive overview of these results, see Sandholm (2010)). Skyrms (1990) has offered a reinterpretation of these evolutionary dynamics in terms of an individual's deliberational process. It will be introduced in this section. In Section 3.3 we will show how this reinterpretation can be extended to DE. We will study two kinds of revision protocols, excess payoff and best response. This makes our investigation fairly representative. If it proves valuable, it may be extended to other kinds of protocols.

Instead of representing evolutionary changes of shares of strategies in a population, the deliberational dynamics suggested by Skyrms (1990) models players' changing beliefs about their and other players' strategies in a normal form game. The initial beliefs are usually incoherent and hence need to be gradually modified according to the deliberational dynamic in question until the incoherence hopefully vanishes and the players converge on beliefs that establish one of the possible NE of a game.[14]

separate these conditional probabilities from i's indecision state in the representation of the first-order beliefs of player i.

[14]The theory of introspective equilibria developed by Kets and Sandroni (2019) in order to account for the phenomenon of homophily rests on a similar idea.

One of the key conceptual differences between evolutionary and deliberational dynamics is the set of conditions used to characterise the stable outcomes of dynamics. In evolutionary models, a population of players in a stable population state is expected to resist perturbations in the numbers of players using different strategies, and hence stable outcomes of evolutionary dynamics are expected to satisfy more stability requirements than are satisfied by each NE (see Sandholm (2010)). In deliberational dynamics, by contrast, NE provide a sufficiently strong stability concept: if players privately revise their beliefs about their own and their opponents' strategies using solely a particular belief revision algorithm, then players' beliefs should never deviate from those prescribed by the dynamic deliberation procedure. Thus, dynamic deliberation offers a novel and conceptually unique justification of NE that relies on players' private revisions of beliefs in a one-agent decision problem.

Skyrms's idea can be formally represented as a model in which a set of $m \geq 2$ players $I := \{1, \ldots, m\}$ are playing a normal form game. Each player $i \in I$ has a finite set of $n_i \geq 2$ pure strategies $S_i := \left\{s_i^1, \ldots, s_i^{n_i}\right\}$ with a typical element s_i^k, and a payoff function $u_i : S \to \mathbb{R}$, where $S := \times_{i \in I} S_i = \left(s^1, \ldots, s^N\right)$ is the set of strategy profiles with a typical element s^h and $N := \prod_{i \in I} n_i$ is the number of strategy profiles in the game. The set of possible opponents' strategy profiles can be defined as $S_{-i} := \times_{j \neq i} S_i$.

Skyrms assumes that the players' deliberational process retains all the independence assumptions of standard game theory, and so the interacting players are represented as believing that their strategy choices are causally and probabilistically independent. The deliberation can thus be modelled as a process of gradual evolution of players' beliefs, where each player i updates her belief system $b_i \in \Delta_{\perp}(S)$, i.e., a probability distribution on the set of strategy profiles S, where $b_i\left(s^h\right) \in [0,1]$ is the probability of profile s^h.

A belief system $b_i \in \Delta_{\perp}(S)$, can be split up into two parts $b_i^i \in \Delta(S_i)$ and $b_i^{-i} \in \Delta_{\perp}(S_{-i})$. On the one hand, b_i^i represents i's probability distribution on her own set of strategies S_i. One may call b_i^i i's indecision state, i's uncertainty about what to do. On the other hand, i's uncertainty about her opponents' strategies can be represented by the distribution b_i^{-i}. Reversely, both together combine into i's belief system b_i.

Moreover, since every $b_i \in \Delta_{\perp}(S)$ assigns a probability distribution on the set of strategies of every player in the set I, we can define i's belief about the profiles of strategies of every player other than $j \in I$ as $b_i^{-j} \in \Delta_{\perp}\left(S_{-j}\right)$, where $b_i^{-j}\left(s_{-j}\right)$ is the probability of j's opponents' profile $s_{-j} \in S_{-j}$, while i's belief about the strategy choices of player j can be defined as $b_i^j \in \Delta\left(S_j\right)$. We assume, however, that b_i^j and b_i^{-j} at the same time represent j's belief system, as i believes it to be. In other words, i assumes her belief system b_i to be mutual knowledge among the players. But this is still her private belief. It is open whether b_i is actually mutual knowledge. (But see below.)

Next, the expected payoff associated with each strategy $s_j^k \in S_j$ of player j relative to player i's belief system b_i is represented with the Lipschitz-continuous payoff function $v_j : \Delta_{\perp}(S) \to \mathbb{R}^{n_j}$ that assigns, to every $b_i \in \Delta_{\perp}(S)$, a vector of expected payoffs $v_j(b_i) := \left(v_j^1(b_i), \ldots, v_j^{n_j}(b_i)\right)$, where $v_j^k(b_i) := \sum_{s_{-j} \in S_{-j}} b_i^{-j}\left(s_{-j}\right) u_j\left(s_j^k, s_{-j}\right)$ is the expected payoff of j's strategy s_j^k relative to i's belief b_i. The indecision state payoff of player j given b_i is defined by the function $\bar{v}_j : \Delta_{\perp}(S) \to \mathbb{R}$ that assigns, to each

$b_i \in \Delta_{\perp}(S)$, an expected payoff $\bar{v}_j(b_i) := \sum_{s_j^k \in S_j} b_i^j\left(s_j^k\right) v_j^k(b_i) = \sum_{s^h \in S} b_i\left(s^h\right) u_j\left(s^h\right)$. Finally, the excess payoff vector for player j given b_i is represented by a function $\hat{v}_j : \Delta_{\perp}(S) \to \mathbb{R}^{n_j}$ that assigns, to each $b_i \in \Delta_{\perp}(S)$, an excess payoff vector $\hat{v}_j(b_i) := \left(\hat{v}_j^1(b_i), \ldots, \hat{v}_j^{n_j}(b_i)\right)$, where $\hat{v}_j^k(b_i) := v_j^k(b_i) - \bar{v}_j(b_i)$ is the excess payoff of s_j^k given b_i.

We can define player i's belief system as forming a NE among players' strategies in a way that is very similar to the definition of a population state establishing a NE in the population.

DEFINITION 9. For any $b_i \in \Delta_{\perp}(S)$, b_i establishes a NE in the private beliefs of player i, iff b_i is such that, for each player $j \in I$ and all $s_j^k \in S_j$,

$$b_i^j\left(s_j^k\right) > 0 \text{ entails } s_j^k \in \arg\max_{s_j^l \in S_j}\left(v_j^l(b_i)\right). \tag{6}$$

Note that so far the interacting players may hold different private belief systems establishing different NE of the game. We will comment on this point below.

However, let us first see how players may reach a NE. One of the deliberational dynamics which guarantees that any convergence of beliefs will establish a NE in player's private beliefs is the excess payoff dynamic — a dynamic generated by an excess payoff revision protocol. In general, a revision protocol is a strategy assignment rule fixing how players choose their strategies in each strategy revision round. An excess payoff revision protocol is a specific strategy revision rule that, in every $b_i \in \Delta_{\perp}(S)$, represents the players as decision-makers who compare the excess payoffs of strategies in b_i and switch to strategies associated with higher excess payoffs than their current strategies. This is how an excess payoff deliberational dynamic based on excess payoff functions $\hat{v}_i$ of some player $i \in I$ is defined:

DEFINITION 10. An excess payoff deliberational revision protocol is a Lipschitz-continuous map $f : \hat{\mathbb{R}}^{n_i} \to \mathbb{R}^{n_i}_{\geq 0}$ that assigns, to every excess payoff vector $\hat{v}_i(b_i) \in \hat{\mathbb{R}}^{n_i}$, a vector of the weights of strategies $f\left(\hat{v}_i(b_i)\right) := \left(f_1\left(\hat{v}_i(b_i)\right), \ldots, f_{n_i}\left(\hat{v}_i(b_i)\right)\right) \in \mathbb{R}^{n_i}_{\geq 0}$, such that $f\left(\hat{v}_i(b_i)\right)' \hat{v}_i > 0$ whenever $\hat{v}_i(b_i) \in \operatorname{int}\left(\hat{\mathbb{R}}^n\right)$, where $\operatorname{int}\left(\hat{\mathbb{R}}^n\right)$ is the interior of $\hat{\mathbb{R}}^n$. $f_k\left(\hat{v}_i(b_i)\right) \in \mathbb{R}_{\geq 0}$ is called the weight of strategy s_i^k.

A desirable feature that an excess payoff protocol should satisfy is that a player should assign a higher weight to a strategy only if it yields a strictly positive excess payoff. This is satisfied by sign-preserving excess payoff revision protocols.

DEFINITION 11. The revision protocol f is sign-preserving iff, for each excess payoff vector $\hat{v}_i(b_i) \in \hat{\mathbb{R}}^n$ and strategy $s_i^k \in S_i$, $f_k\left(\hat{v}_i(b_i)\right) > 0$ iff $\hat{v}_i^k(b_i) > 0$.

The excess payoff deliberational dynamic generated by f represents player i's belief system revision process: in every belief revision period, player i assigns weights on strategies according to protocol f to every distribution $b_i^j \in \Delta\left(S_j\right)$ representing i's beliefs about every player $j \in I$ that is generated by i's belief system b_i. The deliberation dynamic represents a Markovian process, since player i's belief system b_i depends only on the immediately preceding belief system b_i'.

DEFINITION 12. An excess-payoff deliberational dynamic generated by a vector of expected payoff functions $v := (v_1, \ldots, v_m)$ and by a sign-preserving excess payoff revision protocol f is a Lipschitz-continuous map $\beta^{v,f} : \Delta_\perp(S) \to \mathbb{R}^N$ that assigns, to every $b_i \in \Delta_\perp(S)$, an output $\beta^{v,f}(b_i) := \left(\beta^{v_1,f}(b_i), \ldots, \beta^{v_m,f}(b_i)\right) \in \mathbb{R}^N$, where $\beta^{v_j,f}(b_i) := \left(\beta_1^{v_j,f}(b_i), \ldots, \beta_{n_j}^{v_j,f}(b_i)\right) \in \mathbb{R}^{n_j}$ is the vector of the change rates of the probabilities of strategies of player $j \in I$, and where the change rate of the probability of each $s_j^k \in S_j$ is such that

$$\beta_k^{v_j,f}(b_i) := \frac{f_k\big(\hat{v}_j(b_i)\big) - b_i^j\left(s_j^k\right) \sum\limits_{s_j^l \in S_j} f_l\big(\hat{v}_j(b_i)\big)}{\Lambda_i + \sum\limits_{s_j^l \in S_j} f_l\big(\hat{v}_j(b_i)\big)}. \tag{7}$$

We adopt the constant $\Lambda_i \geq 0$ from Skyrms (1990), p. 31. It represents i's index of caution: the higher Λ_i, the slower i increases the probabilities of players' strategies yielding higher than the indecision state expected payoffs.

We may now state the relation between the rest points of the excess payoff deliberational dynamic $\beta^{v,f}$ and the set of belief systems establishing a NE in player i's beliefs.

DEFINITION 13. For any $b_i \in \Delta_\perp(S)$, b_i is a rest point of $\beta^{v,f}$ iff $\beta_k^{v_j,f}(b_i) = 0$ for every $s_j^k \in S_j$ of every $j \in I$.

LEMMA 2. *Let* $\mathrm{bd}\left(\hat{\mathbb{R}}^{n_j}\right)$ *denote the boundary of* $\hat{\mathbb{R}}^{n_j}$. *For any* $b_i \in \Delta_\perp(S)$, $\hat{v}_j(b_i) \in \mathrm{bd}\left(\hat{\mathbb{R}}^{n_j}\right)$ *for every* $j \in I$ *iff* b_i *establishes a NE in the private beliefs of* i.

PROOF. For any $b_i \in \Delta_\perp(S)$ and any $j \in I$, $\hat{v}_j(b_i) \in \mathrm{bd}\left(\hat{\mathbb{R}}^{n_j}\right)$, implies that, for each $s_j^l \in S_j$, $v_j^l(b_i) \leq \bar{v}_j(b_i)$, and so there exists $\varkappa \in \mathbb{R}$, such that, for every $s_j^l \in S_j$, $v_j^l(b_i) \leq \varkappa$ and every $s_j^k \in S_j$ with $b_i^j\left(s_j^k\right) > 0$, $v_j^k(b_i) = \varkappa$, where $\varkappa = \max_{s_j^l \in S_j}\left(v_j^l(b_i)\right)$. Thus, for each $b_i \in \Delta_\perp(S)$, if $\hat{v}_j(b_i) \in \mathrm{bd}\left(\hat{\mathbb{R}}^{n_j}\right)$ for each $j \in I$, then b_i establishes a NE in i's private beliefs. $\square$

LEMMA 3. *For any* $b_i \in \Delta_\perp(S)$, *if* $\hat{v}_j(b_i) \in \mathrm{bd}\left(\hat{\mathbb{R}}^{n_j}\right)$ *for each* $j \in I$, *then* b_i *is a rest point of* $\beta^{v,f}$.

PROOF. For any $b_i \in \Delta_\perp(S)$, such that $\hat{v}_j\left(b_i^{-j}\right) \in \mathrm{bd}\left(\hat{\mathbb{R}}^{n_j}\right)$ for each $j \in I$, sign-preservation of f implies that $\beta_k^{v_j,f}(b_i) = 0$ for each $j \in I$ and $s_j^k \in S_j$. Thus, $\beta_k^{v_j,f}(b_i) = 0$ for each $s_j^k \in S_j$ and $j \in I$ implies that b_i is a rest point of $\beta^{v,f}$. $\square$

PROPOSITION 3. *For any $b_i \in \Delta_\perp(S)$, b_i is a rest point of $\beta^{v,f}$ iff b_i establishes a NE in the private beliefs of i.*

PROOF. Follows immediately from lemmas 2 and 3. □

Let us turn now to our second paradigmatic dynamics, the best response dynamics introduced by Gilboa and Matsui (1992). Under a deliberational interpretation of this dynamic, each player $i \in I$ revises her beliefs in every $b_i \in \Delta_\perp(S)$ by assigning a higher weight to each strategy that is a best response in b_i. If more than one strategy is a best response, the players split the weight equally among all such strategies. The deliberating players are assumed to be cautious and thus only gradually increase the weight of their best responses. As a consequence, strategies that are not best responses only diminish, but do not entirely lose their positive weights.

DEFINITION 14. Let $Br_i(b_i) := \left\{ s_i^k \in S_i : s_i^k \in \arg\max_{s_i^l \in S_i} \left(v_i^l(b_i) \right) \right\} \subseteq S_i$ denote the set of pure best responses of i given b_i. The cautious best response deliberational revision protocol of player i is a map $r_i : \mathbb{R}^{n_i} \to \Delta(S_i)$ that assigns, to vector $v_i(b_i) \in \mathbb{R}^{n_i}$ given b_i, a vector of relative weights of strategies $r_i(v_i(b_i)) := \left(r_i^1(v_i(b_i)), \ldots, r_i^{n_i}(v_i(b_i)) \right) \in \Delta(S_i)$, where $r_i^k(v_i(b_i)) \in [0,1]$ is the relative weight of i's strategy s_i^k, defined by

$$r_i^k(v_i(b_i)) := \begin{cases} \dfrac{b_i^i(s_i^k)}{\sum_{s_i^l \in Br_i(b_i)} b_i^i(s_i^l)} & \text{if } s_i^k \in Br_i(b_i); \\ 0 \text{ otherwise.} \end{cases} \tag{8}$$

DEFINITION 15. A cautious best response deliberational dynamic generated by the vector v of expected payoff functions and a vector of best response protocols $r := (r_1, \ldots, r_m)$ is a map $\beta^{v,r} : \Delta_\perp(S) \to \mathbb{R}^N$ that assigns, to each $b_i \in \Delta_\perp(S)$, an output $\beta^{v,r}(b_i) := \left(\beta^{v_1,r_1}(b_i), \ldots, \beta^{v_m,r_m}(b_i) \right) \in \mathbb{R}^N$, where $\beta^{v_j,r_j}(b_i) := \left(\beta_1^{v_j,r_j}(b_i), \ldots, \beta_{n_j}^{v_j,r_j}(b_i) \right)$ is the vector of the rates of change of probabilities of strategies of $j \in I$ and $\beta_k^{v_j,r_j}(b_i) \in \mathbb{R}$ is the change rate of probability of $s_j^k \in S_j$ defined by

$$\beta_k^{v_j,r_j}(b_i) := \frac{r_j^k(v_j(b_i)) - b_i^j\left(s_j^k\right)}{\Lambda_i + \sum_{s_j^l \in S_j} r_j^l(v_i(b_i))} = \frac{r_j^k(v_j(b_i)) - b_i^j\left(s_j^k\right)}{1 + \Lambda_i}. \tag{9}$$

The relationship between the rest points of the cautious best response deliberational dynamics and the NE of the game is straightforward.

DEFINITION 16. For each $i \in I$ and $b_i \in \Delta_\perp(S)$, b_i is a rest point of $\beta^{v,r}$ iff $\beta_k^{v_j,r_j}(b_i) = 0$ for every $s_j^k \in S_j$ of every $j \in I$.

PROPOSITION 4. *For any $b_i \in \Delta_\perp(S)$, b_i is a rest point of $\beta^{v,r}$ iff b_i establishes a NE in the private beliefs of i.*

PROOF. If b_i is a rest point of $\beta^{v,r}$, then it follows that, for each $j \in I$, $v_j^l(b_i) \leq \varkappa$ for every $s_j^l \in S_j$ and, for every $s_j^k \in S_j$, such that $b_i^j\left(s_j^k\right) > 0$, $v_j^k(b_i) = \varkappa$, where $\varkappa = \max_{s_j^l \in S_j}\left(v_j^l(b_i)\right)$. Thus, for each $b_i \in \Delta_\perp(S)$, if $\beta_k^{v_j,r_j}(b_i) = 0$ for every $\mathrm{s}_j^k \in S_j$ of every $j \in I$, then b_i establishes a NE in i's private beliefs. □

However, so far each player has only private beliefs about herself and the other players. Whether she uses a sign-preserving excess payoff or a cautious best response protocol, her deliberational dynamics generates a solution trajectory that converges to one of the game's NE, her private NE, if it converges at all. So far, the players may be totally out of sync. Of course, if the players' initial uncertainty about themselves, others, and the deliberational dynamics are mutual knowledge, then their deliberation will converge to the same NE. Does this improve upon the mutual knowledge assumptions in Section 3.1? It does not seem so. However, this much mutual knowledge is not required in the present case. If the initial beliefs of the players lead the players to the same NE under some deliberational dynamics, then they will all converge to that NE without any shared knowledge about each others' initial private beliefs.

3.3 *Deliberational Dynamics and Dependency Equilibria*

The previous section treated the standard models of deliberational dynamics. Thereby, we are well prepared to modify the deliberational models so as to accommodate probabilistic entanglement or dependencies in players' beliefs about each other's strategy choices. In this type of model the belief system of each player $i \in I$ can be represented as a probability distribution $p_i \in \Delta(S)$, where $p_i\left(s^h\right) \in [0,1]$ is the probability of profile $s^h \in S$. Note we now drop the independence assumption of the previous section. Again, i's belief system p_i splits into two parts p_i^i and p_i^{-i}. Firstly, i's indecision state concerning her own strategies is represented by the distribution $p_i^i \in \Delta(S_i)$. Secondly, i's beliefs about her opponents' strategy profiles is represented by a conditional distribution p_i^{-i}, where $p_i^{-i}\left(s_{-i}|s_i^k\right)$ represents i's belief about opponents' strategy profile $s_{-i} \in S_{-i}$ conditional on i's strategy s_i^k. Reversely, p_i^i and p_i^{-i} combine to the full distribution p_i.

Given her indecision state, player i has not only conditional, but also unconditional beliefs about the strategies of each player $j \in I$, which are represented by the distribution $p_i^j \in \Delta\left(S_j\right)$. Moreover, i's belief system p_i contains a conditional probability distribution p_i^{-j} on the strategy profiles of all players other than player j conditional on strategies of j. Thus, $p_i^{-j}\left(s_{-j}|s_j^k\right) \in [0,1]$ represents i's probability of j's opponents' profile $s_{-j} \in S_{-j}$ given j's strategy s_j^k. As before, though, we assume that p_i^j and p_i^{-j} constitute j's belief system, as i believes it to be. That is, i again believes her belief system to be mutual knowledge among the players, whether or not this is actually the case.

Relative to player i's belief system p_i, the conditional expected payoffs of player $j \in I$ can be represented by a Lipschitz-continuous expected payoff function $v_j : \Delta(S) \to \mathbb{R}^{n_j}$ that assigns, to every $p_i \in \Delta(S)$, a vector of conditional expected payoffs $v_j(p_i) :=$

$\left(\upsilon_j^1\left(p_1\right),\ldots,\upsilon_j^{n_j}\left(p_i\right)\right)$, where $\upsilon_j^k\left(p_i\right):=\sum_{s_{-j}\in S_{-j}}p\left(s_{-j}|s_j^k\right)u_j\left(s_j^k,s_{-j}\right)$ is the conditional expected payoff of pure strategy $s_j^k\in S_j$ given p_i. The indecision state payoff of j relative to any $p_i\in\Delta\left(S\right)$ is defined as a function $\bar{\upsilon}_j:\Delta\left(S\right)\to\mathbb{R}$ that assigns, to any $p_i\in\Delta\left(S\right)$, an indecision state payoff $\bar{\upsilon}_j\left(p_i\right):=\sum_{s_j^k\in S_j}\left(\upsilon_j^k\left(p_i\right)p_i^j\left(s_j^k\right)\right)=\sum_{s^h\in S}p\left(s^h\right)u_i\left(s^h\right)$. Finally, the excess payoffs of j's strategies are defined via a function $\hat{\upsilon}_j:\Delta\left(S\right)\to\mathbb{R}^{n_j}$ that assigns, to each $p_i\in\Delta\left(S\right)$, a vector of excess payoffs $\hat{\upsilon}_j\left(p_i\right):=\left(\hat{\upsilon}_j^1\left(p_i\right),\ldots,\hat{\upsilon}_j^{n_j}\left(p_i\right)\right)$, where $\hat{\upsilon}_j^k\left(p_i\right):=\upsilon_j^k\left(p_i\right)-\bar{\upsilon}_j\left(p_i\right)$ is the excess payoff of strategy $s_j^k\in S_j$ given p_i.

Having represented the players' entangled belief systems according to the lights of player i, the definition of $p_i\in\Delta\left(S\right)$ establishing a DE in i's private beliefs is structurally identical to the definition of $b_i\in\Delta_\perp\left(S\right)$ establishing a NE in the case where i believes players' strategies to be independent.

DEFINITION 17. For any $p_i\in\Delta\left(S\right)$, p_i establishes a DE in the private beliefs of i iff, for all $j\in I$ and $s_j^l\in S_j$,

$$p_i^j\left(s_j^l|s_i^k\right)>0 \text{ entails } s_j^l\in\arg\max_{s_j^z\in S_j}\left(\upsilon_j^z\left(p_i\right)\right). \tag{10}$$

Now we can generalize our previous results to entangled belief systems and their DE, for both kinds of dynamics considered. First, the excess payoff revision protocol f can be directly applied to the case of probabilistic dependence amongst players' beliefs.

DEFINITION 18. An excess payoff deliberational dynamic generated by a vector of expected payoff functions $\upsilon:=\left(\upsilon_1,\ldots,\upsilon_m\right)$ and by the sign-preserving excess payoff deliberational revision protocol f is a map $\beta^{\upsilon,f}:\Delta\left(S\right)\to\mathbb{R}^N$ that assigns, to any $p_i\in\Delta\left(S\right)$, an output $\beta^{\upsilon,f}\left(p_i\right):=\left(\beta^{\upsilon_1,f}\left(p_i\right),\ldots,\beta^{\upsilon_m,f}\left(p_i\right)\right)$, where $\beta^{\upsilon_j,f}\left(p_i\right):=\left(\beta_1^{\upsilon_j,f}\left(p_i\right),\ldots,\beta_{n_j}^{\upsilon_j,f}\left(p_i\right)\right)$ is the vector of the change rates of probabilities of strategies of player $j\in I$ and where the probability change rate $\beta_l^{\upsilon_j,f}\left(p_i\right)\in\mathbb{R}$ of strategy $s_j^l\in S_j$ is defined as

$$\beta_l^{\upsilon_j,f}\left(p_i\right):=\frac{f_l\left(\hat{\upsilon}_j\left(p_i\right)\right)-p_i^j\left(s_j^l|s_i^k\right)\sum\limits_{s_j^z\in S_j}f_z\left(\hat{\upsilon}_j\left(p_i\right)\right)}{\Lambda_i+\sum\limits_{s_j^z\in S_j}f_z\left(\hat{\upsilon}_j\left(p_i\right)\right)}. \tag{11}$$

DEFINITION 19. For any $p_i\in\Delta\left(S\right)$, p_i is a rest point of $\beta^{\upsilon,f}$ iff $\beta_l^{\upsilon_j,f}\left(p_i\right)=0$ for every $j\in I$ and $s_j^l\in S_j$.

Since the dynamics $\beta^{\upsilon,f}$ and $\beta^{\upsilon,f}$ (Definition 12) are generated by the same protocol f, it is straightforward to check that proofs that are structurally identical to proofs of lemmas 2 and 3 can be used to derive an analogous result concerning the relationship between the rest points of $\beta^{\upsilon,f}$ and the set of dependency equilibria of the game.

LEMMA 4. *For any $p_i \in \Delta(S)$, $\hat{\upsilon}_j(p_i) \in \mathrm{bd}\left(\hat{\mathbb{R}}^{n_j}\right)$ for every $j \in I$ iff p_i establishes a DE in the private beliefs of i.*

PROOF. Structurally identical to proof of Lemma 2. □

LEMMA 5. *For any $p_i \in \Delta(S)$, if $\hat{\upsilon}_j(p_i) \in \mathrm{bd}\left(\hat{\mathbb{R}}^{n_j}\right)$ for every $j \in I$, then p_i is a rest point of $\beta^{\upsilon,f}$.*

PROOF. Structurally identical to proof of Lemma 3. □

THEOREM 3. *For any $p_i \in \Delta(S)$, p_i is a rest point of $\beta^{\upsilon,f}$ iff p_i establishes a DE in i's private beliefs.*

PROOF. Follows immediately from lemmas 4 and 5. □

In the same way, we can transfer the cautious best response dynamic to the case of entangled belief systems.

DEFINITION 20. A cautious best response deliberational dynamic generated by the expected payoff function υ and the cautious best repsonse revision protocol r is a map $\beta^{\upsilon,r} : \Delta(S) \to \mathbb{R}^N$ that assigns, to each $p_i \in \Delta(S)$, an output $\beta^{\upsilon,r}(p_i) := \left(\beta^{\upsilon_1,r_1}(p_i), \ldots, \beta^{\upsilon_m,r_m}(p_i)\right) \in \mathbb{R}^N$, where $\beta^{\upsilon_j,r_j}(p_i) := \left(\beta_1^{\upsilon_j,r_j}(p_i), \ldots, \beta_{n_j}^{\upsilon_j,r_j}(p_i)\right) \in \mathbb{R}^{n_j}$ is the vector of the change rates of probabilities of strategies of $j \in I$ and $\beta_l^{\upsilon_j,r_j} \in \mathbb{R}$ is the change rate of probability of $s_j^l \in S_j$ defined by

$$\beta_l^{\upsilon_j,r_j}(p_i) := \frac{r_j^l(\upsilon_j(p_i)) - p_i^j\left(s_j^l|s_i^k\right)}{\Lambda_i + \sum_{s_j^z \in S_j} r_j^z(\upsilon_i(p_i))} = \frac{r_j^k(\upsilon_j(p_i)) - p_i^j\left(s_j^l|s_i^k\right)}{1 + \Lambda_i}. \tag{12}$$

DEFINITION 21. For any $p_i \in \Delta(S)$ of any $i \in I$, p_i is a rest point of $\beta^{\upsilon,r}$ iff $\beta_l^{\upsilon_j,r_j}(p_i) = 0$ for every $j \in I$ and $s_j^l \in S_j$.

THEOREM 4. *For any $p_i \in \Delta(S)$, p_i is a rest point of $\beta^{\upsilon,r}$ iff p_i establishes a DE in the private beliefs of i.*

PROOF. Structurally identical to proof of Proposition 4. □

Both kinds of dynamics are, however, not ideally adapted to the nature of DE. For, note that neither the excess payoff nor the cautious best response dynamics ensure the preservation of player i's belief that players' strategies are correlated, because both dynamics are based on protocols that shift weights on each players' strategies independently. Thus, i's belief revision procedure may gradually erode player i's initial belief in the probabilistic entanglement of players' strategies. The correlations between players' strategies can, however, be preserved by revision protocols that put positive weights on players' joint strategies that satisfy one of the criteria of fitness represented by the standard revision protocols. This is most easily exemplified with the best response protocol. Hence, let us, as a final consideration, modify the cautious best response protocol to what we call a cautious *joint* best response protocol which represents a player as favoring only strategy profiles that are constituted by best responses of all interacting players.

DEFINITION 22. Let $\upsilon^{p_i} := \left(\upsilon_1\left(p_i\right), \ldots, \upsilon_m\left(p_i\right)\right)$ denote the vector of interacting players' expected payoff vectors relative to the belief system p_i of player i, and let $Br\left(p_i\right) := \left\{s^h \in S : s^h \in \arg\max_{s^z \in S}\left(p_i\left(s^z\right) u_j\left(s^z\right)\right) \text{ for every } j \in I\right\} = \times_{j \in I} Br_j\left(p_i\right)$ denote the set of strategy profiles in which each profile contains only strategies that are best responses of all players in I relative to p_i. The cautious joint best response deliberational revision protocol, then, is a map $\chi_i : \mathbb{R}^N \to \Delta\left(S\right)$ that assigns, to any vector of players' expected payoff vectors υ^{p_i}, a vector of relative weights of strategy profiles $\chi_i\left(\upsilon^{p_i}\right) := \left(\chi_i^1\left(\upsilon^{p_i}\right), \ldots, \chi_i^N\left(\upsilon^{p_i}\right)\right) \in \mathbb{R}^N$, where $\chi_i^h\left(\upsilon^{p_i}\right) \in \mathbb{R}$ is the relative weight of strategy profile $s^h \in S$ defined by

$$\chi_i^h\left(\upsilon^{p_i}\right) := \begin{cases} \dfrac{p_i\left(s^h\right)}{\sum_{s^z \in Br\left(p_i\right)} p_i\left(s^z\right)} \text{ iff } s^h \in Br\left(p_i\right); \\ 0 \text{ otherwise.} \end{cases} \tag{13}$$

The label “joint" marks a distinctive change. The idea is not that a player indiscriminately and proportionally favors all of her best responses (in the sense of maximizing conditional expected utility). The idea is rather that she favors only those strategy profiles in which her best response is combined with best responses of the other players, and she does so proportionally in case there should be several such profiles. In this sense, it is still a best response protocol, but one discriminately favoring profiles of best responses of all players.

This protocol generates the following dynamics:

DEFINITION 23. The cautious joint best response deliberational dynamic generated by the expected payoff function υ and a vector of cautious joint best response revision protocols $\chi := \left(\chi_1, \ldots, \chi_m\right)$ is a map $\beta^{\upsilon,\chi} : \Delta\left(S\right) \to \mathbb{R}^N$ that assigns, to any $p_i \in \Delta\left(S\right)$, an output $\beta^{\upsilon,\chi} := \left(\beta_1^{\upsilon,\chi}\left(p_i\right), \ldots, \beta_N^{\upsilon,\chi}\left(p_i\right)\right) \in \mathbb{R}^N$, where each $\beta_h^{\upsilon,\chi}\left(p_i\right) \in \mathbb{R}$ is the change rate of the probability of strategy profile s^h defined by

$$\beta_h^{\upsilon,\chi}\left(p_i\right) := \frac{\chi_i^h\left(\upsilon^{p_i}\right) - p_i\left(s^h\right)}{1 + \Lambda_i} \tag{14}$$

Again, we have slowed down the responsiveness of the process by adding an a priori weight Λ_i. The simplest way to conceive of such a joint best response deliberational revision protocol is a modified Brown-Robinson process. In the original process for two-person games, each player makes a statistic of the actual or fictitious play of the other player and uses the observed relative frequencies as her expectation of the other's strategy in the next play. Instead, though, we may conceive of the players as making a statistic of the joint actual or fictitious play of strategy profiles and as basing their conditional expectations for the next play on this joint statistic, which will usually display dependencies. Indeed, this seems to be the more reasonable statistic to record. Why be content with a partial statistic when a full one is available? Thus, this modified Brown-Robinson process offers a simple model as to how entangled belief systems might arise.

Once more, we arrive at the desired results:

DEFINITION 24. For any $p_i \in \Delta(S)$, p_i is a rest point of $\beta^{\upsilon,\chi}(p_i)$ iff $\beta_h^{\upsilon,\chi}(p_i) = 0$ for every $s^h \in S$.

THEOREM 5. *For any $p_i \in \Delta(S)$, p_i is a rest point of $\beta^{\upsilon,\chi}(p_i)$ iff p_i establishes a DE in the private beliefs of player $i \in I$.*

PROOF. If $\beta_h^{\upsilon,\chi}(p_i) = 0$ for every $s^h \in S$, it follows that, for every $j \in I$ and every $s^g \in S$, $u_j(s^g)\, p_i(s^g) \leq \kappa$ and, for every $s^h \in S$, such that $p_i\left(s^h\right) > 0$, $u_j\left(s^h\right) p_i\left(s^h\right) = \kappa$, where $\kappa := \max_{s^z \in S} u_j(s^z)\, p_i(s^z)$. Thus, it follows that if $\beta^{\upsilon,\chi}(p_i) = 0$ for every $s^h \in S$, then p_i establishes a DE in the private beliefs of i. $\square$

Theorems 3, 4, and 5 specify the relation between DE and rest points for the three deliberational dynamics for entangled belief systems considered here. Convergence to a rest point is not guaranteed, also not in the case of NE. But, if the deliberational dynamics converges, it must converge to a DE.

4. COMPARISONS

We have indicated in Section 1 that probabilistic dependencies between players' strategies have been previously countenanced by the literature in various ways, resulting in a number of equilibrium notions which are similar to, but not identical with, the novel notion of DE. We need a comparative section to sort this out. At the same time this is intended as a plea for the significance of DE. In our view they deserve a central place in this variegated field.

Of first importance is getting clear on the relation between DE and CE, for the simple reason that the syntactic difference between (2) and (3) consists in just one prime, while we meet quite different constructions in the other comparisons. This is the task of Section 4.1. Afterwards we turn to all the other comparisons.

4.1 *Dependency and Correlated Equilibria*

So far, we have grasped the difference between CE and DE only superficially, although it proved to be dramatic in some of the examples. Still, the significance of the switch from CE to DE may not yet be fully clear. To do better, we should attend to the other and perhaps more important version of CE and ponder whether there is an analogous version of DE. This will also deliver a possible reason for entangled belief systems.

Concerning Definition 2, Osborne and Rubinstein (1994) (p. 47) remark that CE thus defined "may have no natural interpretation", presumably because according to this definition the players' strategies are directly presented as correlated without any indication where the correlation comes from. Therefore they first define CE in the original way of Aumann (1974) and merely derive our Definition 2 as an equivalent characterization. The original idea is that the players make their choices dependent on what goes on in the world, or rather on the private information they receive about the world, which may differ from player to player. Since there is correlation in the world, there is also correlation in the information players receive and thus finally in their choices.

Formally, this goes like this: In addition to a game $\gamma = \left(I, (S_i)_{i\in I}, (u_i)_{i\in I}\right) \in \Gamma$, we assume a (finite) probability space (Ω, μ), where Ω is a set of states and μ a probability measure on Ω. (Ω, μ) is assumed to be common knowledge among the players. The literature usually goes on to assume an information partition $\mathfrak{P}_i$ of Ω for each player i. It is perhaps more perspicuous to directly structure the set Ω of states accordingly. That is, without loss of generality we may assume that a state $\omega \in \Omega$ is a vector $\omega = (\omega_1, \dots, \omega_n)$ of private states, and each player i gets informed about her private state $\omega_i \in \Omega_i$ in the state ω. So, $\Omega = \times_{i=1}^n \Omega_i$, and Ω_i represents the information partition $\mathfrak{P}_i$ of player i.

Now, each player i adopts a strategy σ_i for responding to the information ω_i she receives. That is, a strategy σ_i of player i is a function from Ω_i into S_i. To disambiguate, let's label σ_i an extended strategy, since we have called the elements of S_i strategies as well. Σ_i is the set of all extended strategies of i, and $\Sigma = \times_{i=1}^n \Sigma_i$ is the set of extended strategy profiles, with a typical element $\sigma = (\sigma_i, \sigma_{-i})$. Then we can define:

DEFINITION 25. An extended strategy profile $\sigma \in \Sigma$ is a correlated equilibrium (CE+) of the game γ extended by (Ω, μ) if and only if for each player i, each private state ω_i with $\mu(\omega_i) > 0$, and each extended strategy σ_i' of player i

$$\sum_{\omega\in\Omega} u_i\left(\sigma\left(\omega\right)\right)\mu\left(\omega\right) = \sum_{\omega_{-i}\in\Omega_{-i}} u_i\left(\sigma_i\left(\omega_i\right), \sigma_{-i}\left(\omega_{-i}\right)\right)\mu\left(\omega_{-i}\mid\omega_i\right)$$
$$\geq \sum_{\omega_{-i}\in\Omega_{-i}} u_i\left(\sigma_i'\left(\omega_i\right)\sigma_{-i}\left(\omega_{-i}\right)\right)\mu\left(\omega_{-i}\mid\omega_i\right). \quad (15)$$

That is, after receiving private information ω_i no player i can improve her situation by changing to a different extended strategy.

On the one hand, this is a formally more complicated definition, since this extension (Ω, μ) is part of a CE+. There are infinitely many structurally different CE+ for one and the same game γ. On the other hand, it explains how the correlation of the players' choices come about. There is some artificial or natural random device (Ω, μ) on which the players make themselves dependent in drawing their private information from it. If the random device is suitably chosen, all players can profit from this dependence.

It is a short way from Definition 25 of CE+ back to Definition 2 of canonical CE. We may identify the state space Ω with the set S of strategy profiles. μ thereby turns into a distribution over S itself. There is no additional structure beyond the game γ. The random device thereby turns into a kind of oracle (or mediator) informing each player i about her personal result of the random device, i.e., telling her "choose strategy s_i". Myerson (1991) (p. 253) then assumes that each player still has a choice how to respond to the oracle, i.e., each player i chooses an extended strategy σ_i which is now a function from S_i into S_i (where the first S_i is interpreted as an information space and the second as an option space). Thus, Myerson arrives at the definition that the distribution μ over S and the extended strategy profile σ form a canonical correlated equilibrium if and only if for each player i, each strategy s_i of i with $\mu(s_i) > 0$ and each extended strategy σ'_i of i

$$\sum_{s_{-i}\in S_{-i}} u_i(\sigma_i(s_i), \sigma_{-i}(s_{-i}))\mu(s_{-i}|s_i) \geq \sum_{s_{-i}\in S_{-i}} u_i(\sigma'_i(s_i), \sigma_{-i}(s_{-i}))\mu(s_{-i}|s_i) \qquad (16)$$

Again, no player can improve her situation by changing to a different extended strategy. In a final step, Myerson proves that it is optimal for each player to follow the oracle, i.e., to choose the extended strategy σ_i with $\sigma_i(s) = s_i$. Thus, (16) reduces to our first definition of canonical CE given by (2).

Now, if canonical CE "may have no natural interpretation", this verdict presumably carries over to canonical DE as well. The urgent question hence is: Is there a similar story for canonical DE which provides a derivation of (3) analogous to the one of (2) from (16)? This story again extends a game γ by some state space Ω together with a distribution μ over Ω. But how are we to interpret the state space Ω now? It seems clear that an ordinary state space representing a natural or artificial random device won't do. If the players make their actions dependent on such conditions, this can result only in a CE. Given all the players' private information or even the true state of Ω, the players' choices are independent, and given only player's i private information, i's choice becomes independent from the other players' choices, which remain correlated due to lack of their private information. This is not the idea of DE.

As indicated in Section 1, Spohn (2003, pp. 243ff.) suggested that the players' mental set-ups, their subjective views of their decision situations, may be causally entangled. That is, they may not be merely causally connected by having a common cause in the form of an information structure, of which each player gets her private glimpse. They may be more intimately entangled, say, by having evolved in an interactive history of joint belief and desire formation, e.g., in an evolutionary process, as we have begun to study in Section 3.3. This causal entanglement then results in a probabilistic dependence between the players' mental set-ups. Let us not try to specify this in detail, as Spohn (2003) has at least started. However, capturing this idea in the abstract will lead us to a schematic story underlying canonical DE.

The idea now is that Ω_i represents the set of possible decision situations or mental set-ups of player i and that the state space $\Omega = \times_{i=1}^{n}\Omega_i$ consists of profiles of possible mental set-ups of all players. What do the players choose to do? We can't say because we have not further specified these mental set-ups. However, we may assume that the rules of rationality determine what to do within any mental set-up (including a tiebreak rule, if necessary). As above, we may say that player i has an extended strategy σ_i, i.e., a function from Ω_i into S_i. However, now σ_i is not an object of choice varying across a space Σ_i of possible strategies of player i. It is rather a fixed function from Ω_i to S_i determined by the rules of rationality, which are not specified in our abstract model. As said, we would have to provide the $\omega_i \in \Omega_i$ with structure to specify these rules. Likewise for the other players. So, we should assume a fixed function σ from Ω into S, which determines the outcome of the game according to the mental set-ups of the players. Let us call σ the rational law pertaining to the state space Ω. This may sound strange. However, it is not to say that the players don't have a choice. In each decision situation they do have a choice what to do. It's only that the rational decision is determined by the unspecified rules of rationality. The players are not denied their power to willfully deviate from these rules, and they may be prone to error, but both only to their detriment.[15]

Additionally, we assume that the space Ω of mental profiles is governed by a probability measure μ which represents the social pattern governing the mental set-ups of the players with all its potential dependencies. In fact, in order to cover the general case, we must assume that μ is a Popper measure on Ω. Whose uncertainty does the distribution μ express? The joint one of the players. As usual, we assume that μ is common or mutual knowledge among the players. As discussed in Section 3.2, we may allow that each player i has her subjective assessment μ_i of the situation. Correspondingly, we find the idea of a subjective correlated equilibrium in the literature. But then each player would play her subjective game, with unpredictable outcomes. Here, we assume that the players conceive the game in the same way, and this requires mutual knowledge of μ.

Is this to say that a player is uncertain also about her own mental set-up? This would fly in the face of game theory as it is usually understood. But we need not see it this way. Discussing a similar worry, Aumann (1987) (p. 8) argues that his model rather describes the point of view of an outside observer. Likewise here: the outside observer may indeed be uncertain about all the players. He just observes the social pattern μ. Player i may well grasp in which decision situation $\omega_i \in \Omega_i$ she ends up. It would be odd to say, though, that she thereby receives differential information. It's not like having exclusive access to, say, a weather report. She is *in* the state ω_i, and she is aware of it. Awareness is exclusive, but not ordinary information.

If everything is fixed by the pattern μ and the rational law σ, what else is there to say? Well, we may ask the players whether they are content with the social pattern thus fixed. Player i may have a global discontent. She may wish society to be entirely different from what it is. It is clear, though, that we cannot comment on this global question. Let's only ask whether i can be content with her local part in the pattern. In this pattern, her expected utility is

$$\mathrm{u}_i := \sum_{\omega\in\Omega} u_i(\sigma(\omega))\mu(\omega) = \sum_{\omega_i\in\Omega_i, \omega_{-i}\in\Omega_{-i}} u_i(\sigma_i(\omega_i), \sigma_{-i}(\omega_{-i}))\mu(\omega_{-i}|\omega_i)\mu(\omega_i). \tag{17}$$

She will not be content if she could receive less than u_i with positive probability. Those set-ups ω_i should be avoided and have probability 0. And she will not be content if opportunities to receive more than u_i are excluded and have probability 0. They should be allowed by μ and σ. Such local discontent concerns only the $\mu(\omega_i)$ and must leave the $\mu(\omega_{-i}|\omega_i)$ untouched. That is, potential discontent is measured by the conditional expected utilities of the various $\sigma_i(\omega_i)$. So, to resume, i would have reason to be discontent

[15] Aumann (1987) (p. 8) discusses the same worry with respect to his conception of the state space and concludes that "'freedom of choice' is not an issue".

if and only if any set-up ω_i with less than maximal conditional expected utility would have positive probability. The idea then is that the pattern μ and the rational law σ form a dependency equilibrium (DE+) if all players are locally content in this sense.[16] In a DE+ the pattern is stable. In formal terms:

DEFINITION 26. The Popper measure μ over the state space Ω is a dependency equilibrium (DE+) relative to the rational law σ iff for all players i, for all set-ups $\omega_i \in \Omega_i$ with $\mu(\omega_i) > 0$, and all $\omega_i' \in \Omega_i$

$$\begin{aligned}\sum_{\omega\in\Omega} u_i(\sigma(\omega))\mu(\omega) &= \sum_{\omega_{-i}\in\Omega_{-i}} u_i(\sigma_i(\omega_i), \sigma_{-i}(\omega_{-i}))\mu(\omega_{-i}|\omega_i) \\ &\geq \sum_{\omega_{-i}\in\Omega_{-i}} u_i(\sigma_i(\omega_i'), \sigma_{-i}(\omega_{-i}))\mu(\omega_{-i}|\omega_i').\end{aligned} \tag{18}$$

We may now take the same formal move as above and identify the state space Ω with the set S of strategy profiles. And instead of quantifying over all set-ups $\omega_i \in \Omega_i$ it suffices to quantify over all strategies $s_i \in S_i$. It is obvious that these two moves reduce (6) to (3) and DE+ to canonical DE.

Note that these considerations partially remove the veil of ignorance from the rational law σ. Even though we have not specified the law, it must be consistent with maximizing conditional expected utility relative to the pattern μ.

Aumann (1987) (p. 8) writes about understanding "personal choice as a state variable": "The chief innovation in our model is that it does away with the dichotomy usually perceived between uncertainty about acts of nature and of personal players." If we are right, we better observe the dichotomy; the two kinds of uncertainty work in different ways. Our model does not assume that a player first receives partial information about the other players by observing her own (mental) state and then makes a choice. Rather her (mental) state, which implies her choice, is embedded in a social pattern generating causal and probabilistic dependence amongst the states of the players (represented by a distribution μ over Ω), which may or may not maximize her conditional expected utility.

Note that talking of choice here is ambiguous. The choice may be the chosen action/strategy, or it may be the decision situation in which there is a choice. In the latter sense, personal choice is a state variable, as Aumann and we have assumed here. And the probabilistic dependence between the players' choices in this sense may well be due to their causal entanglement. This entanglement needs to be unfolded in detail, but it is not a causal mystery. And as described, it entails a probabilistic dependence between the players' choices in the first sense without thereby implying any causal dependence between them. This implication would indeed offend a standard assumption of noncooperative game theory. However, it rests upon a fallacy, as explained in Section 1.

Let us finally clarify the relation of our DE+ to the observations of Brandenburger and Friedenberg (2008) about intrinsic correlation. In intrinsic correlation, the players'

[16] We have seen in Section 2.2 that there are also 'bad' DE, with which players should not be satisfied at all. However, then satisfaction is measured by other criteria than simply conditional expected utility.

belief states themselves are correlated and not some external events on which the players depend. This seems to resemble DE+. However, in intrinsic correlation each player still considers herself to be independent from the other players; the correlation is only among the other players' belief states. So, an equilibrium under intrinsic correlation is still a CE+, and the case of intrinsic correlation does not alter our general comparison of CE+ and DE+.

This may suffice as an explanation of the difference between CE+ and DE+. In the rest of the paper, we shall neglect the extension of the underlying game γ by some state space Ω and focus exclusively on canonical DE.

4.2 *On Conventions and Social Norms*

We had emphasized the significance of Theorem 2. Let us slightly expand on this significance by considering social norms. A crucial purpose of social norms is to establish at least a Pareto-optimal social state in circumstances where individual rationality is apparently unable to do so (though, of course, Pareto-optimality as such is still blind to fairness and justice). The main tool for achieving this is to install a system of sanctions. The expectation of sanctions then rationalizes compliance with the norm. Bicchieri (2006) has developed a detailed account of social norms along these lines.

However, this means changing players' utility functions, i.e., the game.[17] We may instead hope to be able to give an account of social norms without changing the game. Indeed, it was Lewis (1969)'s original intention to give an account of conventions solely in terms of mutual expectations (which may entail a practice of sanctions when the expectations are disappointed). However, he did so only in terms of mutual unconditional expectations, and this led him to the idea that conventions are NE in coordination games.

Vanderschraaf (1995) and Gintis (2009) (ch. 7) observed correctly that conventions or norms are rather characterized by a system of mutual conditional expectations: "I follow norm X if and only if you do so, too; and I expect you to reciprocate." Hence, they proposed to use CE instead of NE for explicating conventions or norms. This works nicely in pure coordination games, in which the sole interest of the players is that they all make the same choice (in the relevant sense). Such games have multiple pure NE, and any mixture of them is a CE — and also a DE, as we have seen in discussing Hi-Lo in Section 2.3. These CE, and DE, embody the strongest form of dependence, insofar as the relevant conditional probabilities (of "you do b if I do a") are all either 1 or 0.[18]

However, this move still does not guarantee Pareto-optimality of outcomes outside the realm of coordination games. Sometimes, CE come closer to Pareto-optimality than

---

[17] This is explicit in Bicchieri (2006) (pp. 2f.) who distinguishes conventions, which solve coordination games, and social norms, which solve 'mixed-motive' games and do so by transforming mixed-motive into coordination games. Since DE apply across the board, there is no immediate need here to reproduce this distinction.

[18] The strong dependence seems to disappear in a pure NE. However, this is a deception, created by the fact that in standard probability theory events having probability 1 or 0 are independent from all other events. Instead, the strong dependence present in a proper mixture of pure NE may be preserved in the limit converging to a pure NE, when probabilities are conceived lexicographically. So, perhaps the pure NE embody independence only superficially and in fact hide maximal dependence.

NE. Often they do not, e.g., in PD. In this case, the CE approach to norms and conventions is unhelpful. But we certainly have social norms demanding cooperation in PD and other public good games. Theorem 2 in Section 2.2 does better in this respect by showing that any kind of Pareto-optimal outcome can be reached through mutual conditional expectations.

Section 4.1 indeed suggested that DE are a better model for explaining social norms than CE. It seems odd to say that social norms work in the way of CE+, by providing an information background on which players can correlate or even coordinate. It seems more adequate to conceive of social norms as directly governing the mental set-ups of the players and generating useful dependencies. Hence, we suggest building social norms and conventions on mutual conditional expectations as described by DE or DE+.

4.3 *Binmore's Appeal to Folk Theorems*

Binmore (1998), Section 3.3, still defends the suitability of NE for describing social norms. In Binmore (2010), he explicitly warns not to confuse the NE in artificial laboratory situations with NE in real life. In real life, he says, we are continuously in the midst of playing repeated games. In such contexts, equilibrium behavior is better described by folk theorems about in(de)finitely repeated games, which have very many equilibria. E.g., he reminds us that even the rejection of a proposal in an ultimatum game can thus be understood as equilibrium behavior (Binmore (2010), p. 147).

It is interesting to compare Theorem 144.3 of Osborne and Rubinstein (1994), Section 8.5, their basic folk theorem, and our Theorem 2. They almost agree in their lowest upper bound: Any Pareto-optimal payoff profile above a threshold can be reached by a DE in the single play as well as by a NE in the indefinite repetition (where the threshold is any other DE in the case of DE and the greatest lower bound in the case of NE). They slightly differ in the greatest lower bound: For DE it is the profile of the maximin payoffs of all players, and for NE in the indefinite repetition it is the profile of minimax payoffs of the players. This difference is due to the fact that folk theorems exploit punishment that may go as far as oppressing a player to her minimax payoff, while DE are about what a player can secure for herself, which is her maximin payoff. Finally, there is a great difference in between. There are very many DE between their boundaries, but they don't form a convex set, while every payoff profile in the convex hull of the boundaries for NE can be an NE in the infinite repetition.

Presently, the relevant — and most pleasing — point is only that both accounts reach all Pareto optimal payoff profiles (above a certain threshold). However, Binmore's position has its costs. Folk theorems apply to an enormously idealized scenario, and they fail to hold under more realistic assumptions. Moreover, Binmore represents players as being concerned about their infinite or indefinite future. What their concerns and attitudes are regarding the present or any single play remains unclear.[19] By contrast, DE represent single plays and players' states of mind within them. So, why should we take

[19] For a more substantial criticism of the idea of referring to folk theorems for an explanation of the evolution of cooperation see Bowles and Gintis (2003) (pp. 433f.).

recourse to a highly involved idealization with questionable informative value concerning single plays? Why not try to capture the effect of the long history of repetitions of a game through an entangled belief system, which rationalizes reciprocal and cooperative behavior in the single play directly and without appeal to this idealization? The social steady state may well be captured by a DE.

4.4 *Kantian Equilibria*

Quite a different, but related comparative perspective is suggested by so-called Kantian economics, as proposed by Laffont (1975), who proceeds from the assumption that "we have changed the ethics of the typical agent so that when he maximizes his utility function he assumes that everybody behaves as he does" (p. 433). On this basis, Roemer (2019) has developed a detailed theory of 'Kantian optimization', the core notion of which is that of a 'Kantian equilibrium'. Roemer emphasizes that he thereby proposes to directly change the decision rule of the players. They follow a "quasi-moral norm ... that is motivated by wanting to do the right thing", where "the 'right thing' is defined in large parts by what the others do" (p. 9). And for him, cooperation is precisely such a quasi-moral norm. He does not specify the epistemic make-up of the players, but if one did, one would have to assume the players to be involved in entangled belief systems.

A severe restriction is that Roemer's Kantian equilibrium applies only to (roughly) symmetric games. Roemer specifies ways to loosen it. Still, these ideas do not work for arbitrary utility profiles. By contrast, we have seen that DE can reconstruct these ideas, while maintaining full generality. Still, DE are entirely different in spirit. The label "Kantian" may suggest an orientation toward social optimality, irrespective of a basis in individual rationality.[20] However, this is not the idea of DE. As explained in Section 4.1, DE are grounded in best-response reasoning and thus firmly rooted in individual rationality. If they can rationalize 'Kantian' behavior, all the better.

4.5 *Self-Similarity Reasoning*

Roemer's Kantian equilibria have a normative feel. However, in a more descriptive spirit, they may also be subsumed under what is known as self-similarity reasoning according to which players tend to think that other player(s) will choose in a way similar to their own. E.g., when you and I are playing the PD, I assume that you are similar to me. If I choose to defect, you are likely to do so as well, and likewise for cooperation. To simplify, the choice is only between joint defection and joint cooperation. And if this is the only choice, then it is clear that joint cooperation is the only rational solution.[21]

However, this kind of reasoning is not well reputed in economics. Ross et al. (1977) already studied it under the label "false consensus effect". Rubinstein and Salant (2016)

[20] As we understand it today; for Kant, practical reason was insight into the moral law, in the first place.

[21] This is the mirror principle of Davis (1977), which he proposed for rationalizing cooperation in PD. It works only for symmetric games. Davis' argument reappears in Hofstadter (1983)'s account of superrationality.

have reconfirmed its presence in a sophisticated experiment. It is clear from those papers, though, that the authors think that this way of reasoning is a case for social psychology and not for a theory of rationality. Note, though, that psychology is full of examples where the behavior of people has appeared irrational according to received standards of rationality, but can be construed as rational relative to alternative standards.[22]

Daley and Sadowski (2017) treat the same phenomenon under the label of magical thinking, which refers to a "cognitive error", namely to "the belief that one's action choice influences one's opponent to choose the same action" (p. 913). Thus, they, too, infer causal from probabilistic dependence, a move we have criticized in the introduction. Still, they give an axiomatic equilibrium analysis of the observed behavior, which they argue to be empirically more adequate than other accounts. Due to their quite different set-up there is no straightforward comparison. In any case, their account is developed for PD and then generalized to symmetric two-person games only. It is confusing that the symmetric Pareto-optimum in symmetric games is allegedly supported either by pure (super-)rationality, by moral considerations, or by a cognitive error.

4.6 *Program Equilibria*

We might also think of computer programs as being susceptible to self-similarity reasoning, in particular when they are produced by the same programmer. However, assumptions about what the other players do need not result from guesswork, they may be due to direct information about the decision program they apply. This leads to the idea of a program equilibrium of Tennenholtz (2004). More precisely, the idea is that players do not directly choose strategies in the underlying game. Rather, they play a metagame: Each player chooses a program and submits it to the other players. Here, the program chosen is a function assigning a possibly mixed strategy to all possible submissions from the other players.[23] By submitting a program a player is committed to it. Thus, each program at the same time functions as a kind of commitment device. In this way, each player knows the programs, the minds, as it were, of the other players, and from the choices of the programs a certain possibly mixed strategy profile results.

Which programs should the players choose? The recommendation is to choose a program equilibrium, which is defined as a NE in the metagame (indeed a pure one; mixing programs can only be done by a further program). The NE in the metagame may result in a large variety of strategy profiles of the underlying game. Tennenholtz (2004) (p. 369) proves that any feasible individually rational payoff vector can be reached by a program equilibrium, where such a vector is one that gives each player at least her minimax payoff. (For a generalization see Peters and Szentes (2012).) Thus, as Tennenholtz emphasizes, cooperation in PD and indeed any Pareto optimal strategy profile can be

[22]A famous example is Wason's selection task, initiating a huge discussion within cognitive psychology. In this task, the majority of people made elementary mistakes in deductive logic (by 'affirming the consequent'), while their inferences are perfectly intelligible in terms of Bayesian reasoning. See Evans (2016).

[23]This sounds circular. How can a function apply to itself, or something of the same set-theoretic type, as an argument? Well, such programs must be computable, and then such self-application is feasible after some encoding, e.g., a Gödel numbering, of the possible computable programs.

reached by a program equilibrium. Hence, as Tennenholtz observes (p.370), program equilibria behave in the same way as NE in in(de)finite repetitions of a game regarding lower and upper limits and—as we may observe now—similar to DE. Still, the underlying construction is obviously different. DE are silent on the way entangled beliefs come about and allow any kind of cause. They do not appeal to any metagame construction. Program equilibria, by contrast, are tailored to that specific construction, which deeply involves them in the theory of computable functions.

4.7 *Translucent Equilibria*

As explained, program equilibria assume full knowledge not only of the actions, but of the mental states, i.e. the programs, of the players. This is an unrealistic assumption. Reversely, in game theory, "typically, players are assumed to be opaque, in the sense that a deviation by one player in a normal-form game does not affect the strategies used by other players" (Capraro and Halpern (2019), p. 372). This is unrealistic as well, a point that was our original motivation for introducing DE. Hence, Halpern and Pass (2018) model games with 'translucent' players, which informally means that the players' minds are partially translucent to the other players. In their approach this means in particular that "a player may believe that if he switches from one strategy to another, the fact that he chooses to switch may be visible to the other players" (Capraro and Halpern (2019), p. 372). They cite facial and bodily clues for such transparency. This corresponds to our idea mentioned in the introduction that prior to the players' actions there may be plenty of causal interaction between the players' mental set-ups.

So, the starting point is very similar. The formal implementation, however, is very different. Halpern and Pass (2018) assume, for each player, a structure of counterfactual beliefs about what the other players would do, if she were to switch her strategy. They further assume what they call common counterfactual belief of rationality, i.e., that it is common knowledge that each player i believes the other players would be rational (in the sense of maximizing unconditional expected utility) even if i were to switch her strategy. Very roughly, they go on defining that a strategy profile constitutes a translucent equilibrium iff there exists an appropriate counterfactual structure in which it is known that the players play this profile while having the common counterfactual belief of rationality. And they show that a strategy profile is a translucent equilibrium iff each player's strategy is individually rational (i.e. at least as good as her maximin strategy) within all minimax rationalizable strategy profiles, as they call them.

Our account is much simpler by assuming only the rationality of players in the sense of maximizing conditional expected utilities as well common knowledge of such rationality. Our approaches resemble in motivation and substantially differ in conceptual structure. Still, the extent to which they arrive at similar results may be worth checking.

4.8 *Evidential Reasoning as Heuristics*

A neutral and more general stance is taken by Al-Nowaihi and Dhami (2015). They allow for evidential reasoning, as they call it, represented by so-called social projection functions (SPFs) by which a player assigns probabilities to the other players' actions conditional on her choice of a mixed strategy. These SPFs may take any shape and work for any game in normal form. Al-Nowaihi and Dhami carefully avoid a causal interpretation of these SPFs. Somehow, my choice is just evidence for what the others do. For them, causal reasoning, which they take as standard in game theory, is the special case in which SPFs are constant and players' choices are probabilistically independent from one another. Al-Nowaihi and Dhami subsume their approach under the influential heuristics and biases program in cognitive psychology. They take SPFs as simple heuristics for resolving uncertainty about the other players' actions, which is used instead of complex higher-order mutual best-response reasoning as captured in standard game theory.

Now, each player maximizes conditional expected utility according to her SPF, and these optimal strategies together with the SPFs are defined to form a consistent evidential equilibrium (CEE) (p. 646) precisely if each player's SPF is correct in the sense of predicting the others' optimal strategies given her own optimal strategy. With causal reasoning, i.e., constant SPFs, CEE reduce to NE.

CEE look very much like DE. Among the surprisingly many non-standard equilibrium notions known to us, CEE come closest to DE. In fact, in the case of two-person games the two notions are identical.[24] However, they come apart for ≥ 3-person games. This is so because an SPF expresses expectations about each of the other players individually and thus treats them as independent. DE make no such assumption. For this reason, a collection even of consistent SPFs can never be common knowledge in ≥ 3-person games (unless they are constant), because the probabilistic dependence one player sees between her and the others' strategies cannot be replicated by the other players. By contrast, DE are entailed by such common knowledge, as shown in Section 4.1. This difference reflects the difference in interpretation. Still, the rich examples and comparisons with which Al-Nowaihi and Dhami intend to display the explanatory force of CEE may carry over to DE; indeed, they fully carry over for 2-person games.

4.9 *Evidential Decision Theory*

Philosophers have much less difficulty in seeing others' actions or other states of the world as probabilistically, though not causally, dependent on one's own actions. This is so because since Nozick (1969) philosophers have intensely discussed Newcomb's problem (NP) and the ensuing and still hotly debated distinction between causal and evidential decision theory (CDT and EDT), a distinction that had little impact within economics.

In NP, you have the choice between taking an opaque box containing either nothing or a million dollars and taking this box together with a second containing a thousand dollars. The problem posits that the opaque box has been filled long before your choice by a very reliable predictor. If he predicted that you would take only the opaque box, he has filled it with the million, and if he predicted that you would take both boxes, he has left it empty. CDT recommends that you choose the dominant action of taking both boxes because your choice can't influence a temporally prior prediction. Following

[24] Apart from the technical difference that DE use probabilities conditional on pure strategies, while SPFs assume probabilities conditional on mixed strategies, and these probabilities need not be the corresponding mixtures of the probabilities conditional on pure strategies.

dominance or causal reasoning, you are likely to end up with a thousand dollars only. EDT, by contrast, recommends taking only the opaque box, because doing so is at least evidence for the predictor having put the million into it and thus maximizes conditional expected utility. Following evidential reasoning, you are likely to end up with a million dollars.[25] Maybe CDT is not so reasonable after all.[26]

If such evidential relations hold only in fantastic cases like NP, then one might well set them aside. However, our various points about social norms, self-similarity considerations etc. suggested that such evidential relations are a perfectly mundane phenomenon easily arising in the social world between human agents with a common background.[27] Constantly, we are mutual predictors of one another, not very reliable ones perhaps, but certainly better than random ones, and thereby generate entangled belief systems. One might say that the only rational way of predicting rational people is via standard game theory. Clearly, though, this would beg the question. We better accept such entangled belief systems as a relevant fact and study their rational consequences.

Just as CDT pairs with game theory based on NE, game theory based on DE naturally corresponds with EDT. However, this should only be taken as an offer to the philosophical adherents of EDT. We hope to satisfy the adherents of CDT and standard game theorists as well, namely by providing the causal explanation of entangled belief systems laid out in in Sections 1 and 4.1. We need not take a stance on the debate between CDT and EDT in this paper.

4.10 *Summary*

To wrap up: Game theorists have offered a surprising variety of accounts, matched with nonstandard equilibrium notions, that try to come to terms, descriptively or normatively, with various phenomena which are difficult to account for within standard game theory. It seems to us that DE should take a central place in this field. More detailed comparisons would certainly be desirable, but space is lacking. Let us only add: As we have mentioned, the various accounts offer variegated empirical and experimental evidence. DE were not considered in these experiments. It seems that often such experiments are unable to discriminate between the hypothesis actually confirmed and a similar hypothesis couched in terms of DE. That is, existing evidence for these alternative equilibrium notions could also be interpreted in favor of DE. Again though, it would go much too far to start a detailed analysis of all the relevant material in order to gain a comprehensive assessment of the empirical standing of DE.

[25] According to a survey from 2020, 39% of 1071 philosophers accepted or leaned towards taking two boxes in NP and 31% accepted or leaned towards one-boxing. It is unlikely that the latter are simply confused. See: https://survey2020.philpeople.org/survey/results/4886. In the end, though, the philosophical discussion is more complicated. Quite a few philosophers challenge the common identification of EDT with one-boxing and CDT with two-boxing. See, e.g., Spohn (2012).

[26] However, after Jeffrey (1965/1983) developed EDT and challenged CDT, CDT required more explicit statements, as delivered, e.g., by Gibbard and Harper (1978), Skyrms (1980) (part IIC), and Lewis (1981). The real challenge of the issue is, of course, to get clear about the relation between causation and probability in the context of decision theory, which is not fully explicit in Savage (1954), denied by Jeffrey (1965/1983), and contested since. As mentioned, the leading paradigm here is the interventionist theory of causation, as presented, e.g. in Pearl (2009).

[27] Already Brams (1975) has observed that PD is just a symmetricized version of NP.

5. Conclusion

Our goal has been to present a conceptual generalization of NE to DE and to vindicate the significance and plausibility of this move. Our main accomplishment in this regard has been to show that major branches of game theory such as epistemic game theory and evolutionary game theory (in its deliberational reinterpretation) can just as well be developed on the basis of the concept of DE, yielding intriguing consequences that often deviate from standard game theory and merit further exploration. We thus hope to have laid the foundations for a novel and worthwhile research program.

Appendix

In Section 2.2 we claimed that the limit definition 4 and the lexicographic definition 7 of a DE are essentially equivalent. In order to explain this, we first have to clarify the relation between sequences of probability distributions in $\Delta(S)$ and Popper measures. This is stated in:

Lemma 6. *Let the sequence $(p_r)_{r\in\mathbb{N}}$ of distributions $p_r \in \Delta(S)$ be such that for some $T \subseteq S$ $\lim_{r\to\infty} p_r(A|B)$ exists for all $A, B \subseteq S$ with $B \cap T \neq \emptyset$. Then P defined by $P(A|B) = \lim_{r\to\infty} p_r(A|B)$ for all such $A, B \subseteq S$ is a Popper measure for S. Conversely, for each Popper measure P for S there is a sequence $(p_r)_{r\in\mathbb{N}}$ of distributions $p_r \in \Delta(S)$ and $T \subseteq S$ such that $P(A|B) = \lim_{r\to\infty} p_r(A|B)$ for all $A, B \subseteq S$ with $B \cap T \neq \emptyset$.*

Proof. The first claim is trivial, because each p_r has the properties characteristic of a Popper measure, which are preserved in the limit. Regarding T, observe the last assertion of Lemma 1. Reversely, let $\lambda = (\lambda_0, \ldots, \lambda_q)$ be the lexicographic probability corresponding to P. For each $\epsilon > 0$, let $p'_\epsilon = \sum_{j=1}^{q} \epsilon^j \lambda_j$, and let p_ϵ be the normalization of p'_ϵ. Then, $P(A|B) = \lim_{\epsilon\to 0} p_\epsilon(A|B)$. $\square$

Now we can state the essential equivalence between lexDE and lim DE:

Theorem 6. *If P is a lexDE, then $p = P(.|S)$ is a limDE. Reversely, if p is a limDE, then there is a $P \in \Delta^*_{Popp}(S)$ such that $p = P(.|S)$ and P is a lexDE.*

Proof. Let P be a lexDE. Then as shown in Lemma 2, there is a sequence $(p_r)_{r\in\mathbb{N}}$ of distributions $p_r \in \Delta(S)$ such that $P(A|B) = \lim_{r\to\infty} p_r(A|B)$ for all A, B for which $P(A|B)$ is defined. Let $p = \lim_{r\to\infty} p_r$. Then p is a limDE, because

$$\lim_{r\to\infty} \sum_{s\in S} u_i(s) p_r(s_{-i}|s_i) p_r(s_i) = \sum_{s\in S} u_i(s) P(s_{-i}|s_i) P(s_i) \text{ and} \tag{19}$$

$$\text{for all } s'_i \in S_i \lim_{r\to\infty} \sum_{s_{-i}\in S_{-i}} u_i\left(s'_i, s_{-i}\right) p_r\left(s_{-i}|s'_i\right) = \sum_{s_{-i}\in S_{-i}} u_i\left(s'_i, s_{-i}\right) P\left(s_{-i}|s'_i\right). \tag{20}$$

Reversely, assume that p is a limDE. Thus, there is a sequence $(p_r)_{r\in\mathbb{N}}$ with $p_r \in \Delta^+(S)$ that converges to p, as required. Without loss of generality, we may assume that also

$\lim_{r\to\infty} p_r(A|B)$ exists for all A and B with $B \cap T \neq \emptyset$ for some $T \subseteq S$ for which $T \subseteq \{s_i\} \times S_{-i})$ for all $s_i \in S_i$ and $i \in I$. Why? This is not guaranteed in general. However, if the limit $\lim_{r\to\infty} p_r(A|B)$ does not exist, a limit inferior and a limit superior must exist in any case, and hence we find a subsequence $(p_r)_{r\in M}$ of $(p_r)_{r\in\mathbb{N}}$ for some an infinite subset M of $\mathbb{N}$) for which $\lim_{r\to\infty} p_r(A|B)$ exists. We may repeat the procedure finitely many times for all A and B concerned. Thus, for some infinite $M \subseteq \mathbb{N}$ there is a subsequence $(p_r)_{r\in M}$ of $(p_r)_{r\in\mathbb{N}}$ for which $\lim_{r\to\infty} q_r(A|B)$ exists for all A and $B \neq \emptyset$. Hence, we may assume that $(p_r)_{r\in\mathbb{N}}$ itself is such a suitable sequence. According to Lemma 2, it converges to a Popper measure P. Indeed, $p \in \Delta^*_{Popp}(S)$, due to the assumption about T. And this P must be a lexDE, because equations (19) and (20) apply again. □

Note that there may be many different lexDE corresponding to one limDE, but not vice versa. Note also that not each sequence $(p_r)_{r\in\mathbb{N}}$ converging to a limDE converges to a Popper measure that is a lexDE.

References

Al-Nowaihi, Ali and Sanjit Dhami (2015), "Evidential equilibria: Heuristics and biases in static games of complete information." *Games*, 6, 637–676. [4, 35]

Aumann, Robert (1974), "Subjectivity and correlation in randomized strategies." *Journal of Mathematical Economics*, 1, 67–96. [2, 6, 7, 15, 27]

Aumann, Robert (1987), "Correlated equilibrium as an expression of bayesian rationality." *Econometrica*, 55, 1–18. [6, 29, 30]

Bernheim, B. Douglas (1984), "Rationalizable strategic behavior." *Econometrica*, 52(4), 1007–1028. [16]

Bicchieri, Christina (2006), *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press, Cambridge. [31]

Binmore, Ken (1998), *Game Theory and the Social Contract II: Just Playing*. MIT Press, Cambridge, MA. [32]

Binmore, Ken (2010), "Social norms or social preferences?" *Mind and Society*, 9, 139–157. [4, 32]

Bowles, Samuel and Herbert Gintis (2003), "The origins of human cooperation." In *Genetic and Cultural Evolution of Cooperation* (P. Hammerstein, ed.), 429–443, MIT Press. [32]

Brams, Steven (1975), "Necomb's problem and prisoners' dilemma." *The Journal of Conflict Resolution*, 19, 596–612. [37]

Brandenburger, Adam and Amanda Friedenberg (2008), "Intrinsic correlation in games." *Journal of Economic Theory*, 141, 28–67. [30]

Capraro, Valerio and Joseph Y. Halpern (2019), "Translucent players: Explaining cooperative behavior in social dilemmas." *Rationality and Society*, 31, 371–408. [35]

Cartwright, Nancy (2007), *Hunting Causes and Using Them. Approaches in Philosophy and Economics*. Cambridge University Press, Cambridge. [3]

Daley, Brendan and Philipp Sadowski (2017), "Magical thinking: A representation result." *Theoretical Economics*, 12, 909–956. [34]

Datta, Ruchira S. (2003), "Universality of nash equilibria." *Mathematics of Operations Research*, 28(3), 424–432. [9]

Davis, Lawrence (1977), "Prisoners, paradox, and rationality." *American Philosophical Quarterly*, 14(4), 319–27. [33]

Evans, Jonathan S.B.T. (2016), "A brief history of the wason selection task." In *The Thinking Mind: The Use of Thinking in Everyday Life* (N. Galbraith, ed.), 1–14, Psychology Press, Hove, England. [34]

Fishburn, Peter (1964), *Decision and Value Theory*. Wiley, New York. [2]

Gauthier, David (1986), *Morals by Agreement*. Oxford University Press, Oxford. [14]

Gibbard, Allan and William Harper (1978), "Counterfactuals and two kinds of expected utility." In *Foundations and Applications of Decision Theory* (C. Hooker, J. Leach, and E. McClennen, eds.), 125–162, Dordrecht, Reidel. [37]

Gilboa, Itzhak and Akihiko Matsui (1992), "A model of random matching." *Journal of Mathematical Economics*, 21, 185–187. [21]

Gintis, Herbert (2009), *The Bounds of Reason. Game Theory and the Unification of the Behavioral Sciences*. Princeton University Press, Princeton. [4, 31]

Halpern, Joseph Y. and Rafael Pass (2018), "Game theory with translucent players." *International Journal of Game Theory*, 47, 949–976. [4, 35]

Hofstadter, Douglas (1983), "Dilemmas for superrational thinkers, leading up to a luring lottery." *Scientific American*. [33]

Jeffrey, Richard (1965/1983), *The Logic of Decision*. University of Chicago Press, Chicago. [37]

Karpus, Jurgis and Mantas Radzvilas (2018), "Team reasoning and a measure of mutual advantage in games." *Economics and Philosophy*, 34(1). [14]

Kets, Willemien and Alvaro Sandroni (2019), "A belief-based theory of homophily." *Games and Economic Behavior*, 115, 410–435. [17]

Laffont, Jean-Jacques (1975), "Macroeconomic constraints, economic efficiency, and ethics: An introduction into kantian economics." *Economica*, 42, 430–437. [33]

Lewis, David (1969), *Convention*. Oxford University Press, Oxford. [15, 31]

Lewis, David (1981), "Causal decision theory." *Australasian Journal of Philosophy*, 59, 5–30. [37]

Myerson, Roger B. (1991), *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA. [6, 27]

Nozick, Robert (1969), "Newcomb's problem and two principles of choice." In *Essays in Honor of Carl G. Hempel* (Nicholas Rescher, ed.), 107–33, Dordrecht, Reidel. [36]

Osborne, Martin J. and Ariel Rubinstein (1994), *A Course in Game Theory*. MIT Press, Cambridge, MA. [26, 32]

Pearce, David G. (1984), "Rationalizable strategic behavior and the problem of perfection." *Econometrica*, 52(4), 1029–1050. [16]

Pearl, Judea (2009), *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge. [3, 37]

Peters, Michael and Balázs Szentes (2012), "Definable and contractible contracts." *Econometrica*, 80, 363–411. [34]

Popper, Karl R. (1938), "A set of independent axioms for probability." *Mind*, 47(186), 275–277. [8]

Portakal, Irem and Bernd Sturmfels (2022), "Geometry of dependency equilibria." *arXiv*. [9]

Reichenbach, Hans (1956), *The Direction of Time*. University of California Press, Los Angeles. [2]

Roemer, John (2019), *How We Cooperate: A Theory of Kantian Optimization*. Yale University Press, New Haven. [4, 33]

Ross, Lee, David Greene, and Pamela House (1977), "The "false consensus effect": An egocentric bias in social perception and attribution processes." *Journal of Experimental Social Psychology*, 13(3), 279–301. [33]

Rubinstein, Ariel and Yuval Salant (2016), ""isn't everyone like me?": On the presence of self-similarity in strategic interactions." *Judgment and Decision Making*, 11(2), 168–173. [33]

Sandholm, William H. (2010), *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, Massachusetts. [4, 17, 18]

Savage, Leonard (1954), *The Foundations of Statistics*. Dover, New York. [2, 37]

Skyrms, Brian (1980), *Causal Necessity*. Yale University Press, New Haven. [37]

Skyrms, Brian (1990), *The Dynamics of Rational Deliberation*. Harvard University Press, Cambridge, MA. [4, 17, 20]

Spohn, Wolfgang (1982), "How to make sense of game theory." In *Philosophy of Economics* (W. Stegmüller, W. Balzer, and W. Spohn, eds.), 239–270, Springer, Berlin. [16]

Spohn, Wolfgang (1986), "The representation of popper measures." *Topoi*, 5, 69–74. [8]

Spohn, Wolfgang (2003), "Dependency equilibria and the causal structure of decision and game situations." *Homo Oeconomicus*, 20, 195–255. [3, 4, 6, 7, 11, 14]

Spohn, Wolfgang (2012), "Reversing 30 years of discussion: why causal decision theorists should one-box." *Synthese*, 187, 95–122. [37]

Sugden, Robert (1993), “Thinking as a team: Towards an explanation of nonselfish behaviors.” *Social Philosophy and Policy*, 10, 69–89. [12, 14]

Sugden, Robert (2011), “Mutual advantage, conventions and team reasoning.” *International Review of Economics*, 58(1), 9–20, URL https://doi.org/10.1007/s12232-011-0114-0. [14]

Tennenholtz, Moshe (2004), “Program equilibrium.” *Games and Economic Behavior*, 49, 363–373. [4, 34]

van Fraassen, Bas (1976), “Representational of conditional probabilities.” *Journal of Philosophical Logic*, 5, 417–430. [8]

Vanderschraaf, Peter (1995), “Convention as correlated equilibrium.” *Erkenntnis*, 42, 65–87. [4, 31]