MPRA
Munich Personal RePEc Archive

Rule Rationality

Heller, Yuval and Winter, Eyal

University of Oxford, The Hebrew University of Jerusalem

31 July 2013

Online at https://mpra.ub.uni-muenchen.de/48746/
MPRA Paper No. 48746, posted 31 Jul 2013 15:07 UTC

Rule Rationality

Yuval Heller* Eyal Winter†

July 30, 2013

Abstract

We study the strategic advantages of following rules of thumb that bundle different games together (called *rule rationality*) when this may be observed by one's opponent. We present a model in which the strategic environment determines which kind of rule rationality is adopted by the players. We apply the model to characterize the induced rules and outcomes in various interesting environments. Finally, we show the close relations between act rationality and "Stackelberg stability" (no player can earn from playing first). JEL classification: C72, D82.

1 Introduction

Act rationality is the notion that perfectly rational agents choose actions that maximize their utility in every specific situation. However, both introspection and experimental evidence suggest that people behave differently and follow "rule rationality": they adopt rules of thumb, or modes of behavior, that maximize some measure of average utility taken over the set of decision situations to which that rule applies; then, when making a decision, they choose an action that accords with the rule they have adopted. As a result of bundling together many decision situations, this rule of thumb induces smaller cognitive costs and fewer informational requirements (see Baumol & Quandt, 1964; Harsanyi, 1977; Ellison & Fudenberg, 1993; Aumann, 2008, and the references there).

Reducing the cognitive costs is arguably the main advantage of such rules of thumb, but it is not the only advantage. In this paper we focus on an aspect of rule rationality that has not been studied in the literature. If players tend to bundle games together while interacting with others who gradually learn the structure of bundling, then bundling can serve as a commitment device.[1] *The contribution of this paper is threefold: (1) We present a tractable model in which the properties of the strategic environment determine which kind of rule rationality is adopted by the players; (2) we characterize the induced rules and outcomes in various interesting environments; and (3) we show the close relations between act rationality and "Stackelberg stability" (no player can earn from playing before his opponent).*

A common manifestation of bundling as a commitment device is the tendency to stress moral or ideological principles for strategic reasons. When the US declares that it will never abandon its allies, or

*Department of Economics and Nuffield College, University of Oxford. yuval.heller@economics.ox.ac.uk. URL: https://sites.google.com/site/yuval26/.

†Center for the Study of Rationality and Department of Economics, Hebrew University of Jerusalem. mseyal@mscc.huji.ac.il . URL: http://www.ma.huji.ac.il/∼mseyal/. The author is grateful to the German Israeli Foundation for Scientific Research and to Google for their financial support.

[1] This "commitment" advantage of rule rationality might be of importance also in single-player games as a "self commitment" device with a multiple-selves interpretation.

when a businessman claims that he has never reneged on his promises on moral grounds, they are doing precisely this. Since these commitments refer not to a single game but to a class of games, they are distinguishable from those often discussed in the economics or the political science literature (see, e.g., the seminal work of Schelling, 1960, and the recent papers of Bade et al., 2009; Renou, 2009). Moral principles by their nature bundle multiple situations by means of the action that the principle dictates.

Indeed, the extent to which the commitment to act according to a moral or ideological principle is credible depends on comparing the cost of adhering to the principle relative to the cost of losing the commitment benefits in the future. These are the kind of considerations we employ in our model to define rule rationality. The declaration of such principles can be viewed as a signal that facilitates the gradual learning about one's bundling strategy by others. To be sure, these signals are not fully revealing but the fact that these declarations are very often made and often taken seriously suggests that stressing ideological or moral principles is not merely cheap talk either.

Another form of rule rationality relevant to our context includes firms or organizations that adopt rigid policies applying to a broad class of circumstances without elaborating the rationale of each policy. If a craftsman rejects an offer to accept 50% of his proposed price for doing 50% of the job on the basis the it is "unprofessional" or when a chain store rejects a returned item after the refund deadline even under exceptional circumstances by simply arguing that "it is against our policy", they act within this form of rule rationality.

We now briefly present our model. An *environment* is a finite set of two-player normal-form games that share the same set of feasible actions.[2] The environment is endowed with a function that determines the probability of each game being played. Each environment induces the following two-stage *meta-game:* at stage 1, each player chooses a partition over the set of games; at stage 2, each player chooses an action after observing the element of his partition that includes the realized game and his opponent's partition.[3] A *rule-rational equilibrium* of the environment is defined as a subgame-perfect equilibrium of the meta-game. We interpret such an equilibrium as a stable outcome of a dynamic process of social learning, in which the choices of partitions in the first stage (interpreted as rules of thumb that bundle games together) and the behaviors in the second stage (given the partitions) evolve. A rule-rational equilibrium is *act-rational* if players play a Nash equilibrium in all games, and it is *non-act-rational* if at least one of the players do not best reply in at least one game.

The following motivating example demonstrates that a non-act-rational equilibrium can strictly Pareto dominate the set of of all act rational equilibrium payoff.

Example 1. Consider an environment with the two games shown in Table 1: The Prisoner's Dilemma (G_1) and the Stag-Hunt Game (G_2).

Table 1: Prisoner's Dilemma and Stag Hunt

G_1 (Prisoner's Dilemma)		
	C	D
C	9,9	0,9.5
D	9.5,0	1,1

G_2 (Stag Hunt)		
	C	D
C	9,9	0,8.5
D	8.5,0	1,1

These two games are relatively "similar" in the sense that the maximal difference between the two payoff matrices is 1. Note that in the Prisoner's Dilemma, (D,D) is the unique Nash equilibrium, while the Stag-Hunt game admits two pure Nash equilibria: (D,D) and the Pareto-dominant (C,C). In any

[2]Assuming a fixed set of action is without loss of generality; see Footnotes 6 and 9.
[3]In Section 7.2 we extend our results to a setup with partial observability of the opponent's partition.

act-rational equilibrium of this environment, players play the unique Nash equilibrium (D, D) in ω_1. Thus each player can achieve an expected (ex-ante) payoff of at most 5 in any act-rational equilibrium. Consider the following non-act-rational equilibrium. Both players choose the coarse partition (i.e., bundle together both games), play C on the equilibrium path, and play D (in both games) if any player deviated while choosing his partition. This equilibrium yields both players a payoff of 9. Note that the coarse partition in this example can be interpreted as if each player did not pay attention to small differences in the payoffs (say, up to one utility point).

Section 4 presents additional motivating examples. Example 2 shows an environment in which players choose a coarse partition that pays attention only to the sum of payoffs, and hence they implement an outcome that strictly Pareto-dominates the set of act-rational payoffs in all games (and not only from the *ex-ante* perspective). Example 3 presents an environment that admits only non-act-rational equilibria. Finally, Examples 4 and 5 demonstrate that non-act-rational equilibria can fit experimentally observed behaviors in the "Ultimatum game" and the "Chainstore paradox" (see Camerer, 2003, and the references there).

In Section 5 we characterize rule-rational equilibria in various families of environments. An environment is *redundant* if all its games are strategically equivalent (Moulin & Vial, 1978). Note that such games share the same set of correlated equilibria (Aumann, 1974). Our first results (Propositions 1-4) show that rule rationality in redundant environments can only induce behavior that is consistent with the correlated equilibria of the underlying games; moreover, if the redundant environment is large enough, and each game has a small probability, then the correlated equilibrium of each underlying game can be approximated by a rule-rational equilibrium.

Next, we study environments in which in each game the sum of the payoffs is constant. We show that such constant-sum environments always admit an act-rational equilibrium. They may admit also non-act-rational equilibria, but these equilibria are limited in two ways: they must yield the same ex-ante expected payoff as the act-rational equilibrium, and they are not *robust,* i.e., they are not stable to a perturbation that yields an arbitrarily small probability that the choice of partition is not observed by the opponent. Examples 8-9 show that these properties need not hold if the games are strategically zero-sum (Moulin & Vial, 1978).

Proposition 5 deals with environments in which each game admits a Pareto-dominant action profile, and we show that such environments admit an act-rational equilibrium that strictly Pareto-dominates all non-act-rational equilibria.[4] Next we characterize (Proposition 1) a "folk theorem" result for environments in which a specific action profile is a Nash equilibrium in all games. This is followed by showing that if a specific action is dominant for one of the players in all games in the environment, then the environment admits only act-rational equilibria. Finally, we generalize Example 1 and show that if an environment assigns high enough probability to game g^* that admits strict equilibrium a^* that dominates another strict equilibrium, then there is a rule-rational equilibrium that bundles together all games that have payoff matrices close enough to g^*, and players play a^* in all these games.

In Section 6 we show the close relationship between the existence of act-rational equilibria and of Stackelberg stability. A Nash equilibrium in a game is *Stackelberg-stable* if it yields both players at least their *mixed-action Stackelberg payoff* (Mailath & Samuelson (2006)); that is, no player can guarantee a higher payoff by being a *Stackelberg leader*, i.e., playing first, and having his chosen strategy observed by his opponent before playing. A game is Stackelberg-stable if it admits a Stackelberg-stable equilibrium. Examples of Stackelberg-stable games are constant-sum games, games with a Pareto-dominant action

[4] A similar result was obtained in the framework of value of information in Bassan et al. (2003).

profile, and games in which both players have dominant actions. We first show that if all games in the environment are Stackelberg-stable, then the environment admits an act-rational equilibrium. Second, we show that for any Stackelberg-unstable game G^*, there is an environment in which only G^* is Stackelberg-unstable, and the environment does not admit an act-rational equilibrium. Finally, we show that any act-rational equilibrium of a generic environment is preserved when a Stackelberg-stable game is added to the environment with small enough probability.

We conclude the paper with a discussion of a few extensions and implications of our model (Section 7): (1) we use our model to study the value of information (see the related literature in the following section) and its relations to Stackelberg stability; (2) we extend the model to a setup of partial observability; (3) we discuss the empirical predictions of our model; and (4) we present ideas for future research.

2 Related Literature

This paper is related to and inspired by three strands of literature (in addition to the literature commitments mentioned earlier). The first strand presents an evolutionary foundation for seemingly irrational behavior in strategic environments. Pioneered by Guth & Yaari (1992), the literature on the evolution of preferences shows that if the preferences of an agent are observed by an opponent, then his subjective preferences (which the agent maximizes) may differ from his material preferences (which determine his fitness); see, e.g., Dekel et al. (2007); Winter et al. (2010) interprets these subjective preferences to be "rational emotions".[5] This literature describes stable outcomes of an evolutionary process in which subjective preferences evolve slowly, and behavior (given the subjective preferences) evolves at a faster pace. Recently, a few papers have used a similar methodology to study the evolution of cognitive biases, such as the "endowment effect" (Frenkel et al., 2012) and bounded forward-looking (Heller, 2013). In this paper we adopt a similar approach, and study evolutionary foundations for ignoring the differences between different games, and bundling them together in a single equivalence set.

The second strand of literature studies *categorization.* It is commonly accepted in the psychology literature that people categorize the world around them in order to help the cognitively limited categorizer to deal with the huge amount of information he obtains from his complex environment. Samuelson (2001) presents a model of cognitively limited agents who bundle together different kinds of bargaining games in order to free up cognitive resources for competitive tournament games. Jehiel (2005) models players who bundle together different nodes in an extensive-form game, and best reply to aggregate opponent's behavior in this "analogy set." Azrieli (2009) studies games with many players, and shows that categorizing the opponents into a few groups can lead to efficient outcomes. Mohlin (2011) studies the optimal categorization in prediction problems. Finally, Mengel (2012) studies a learning process in which agents jointly learn which games to bundle together, and how to play in each partition element. The main novelty of this paper relative to this existing literature is the study of the the strategic "commitment" advantages of "observable" categorization, and the presentation of an equilibrium concept for classes of games (rather than for specific games).

The third strand deals with the *value of information.* In a seminal paper, Blackwell (1953) shows that for single-agent decision problems, more information for the agent is always better. This result can also be extended to constant-sum games, but it is not true, in general, for non-constant-sum games. A classical example of this is given in Akerlof (1970), which shows that providing the seller with private information might render any trade impossible, and thereby reduce the welfare of both the seller and the

[5] Also related are the papers that study environments in which a player can choose a delegate (with different incentives) to play on his behalf (see, e.g., Fershtman et al. (1991)).

buyer; see additional examples in Kamien et al. (1990). A few papers characterize specific environments in which information has positive value. Bassan et al. (2003) show that if all games have common interests, then information has a (Pareto-)positive value: more information yields a Pareto-dominant equilibrium (see also Lehrer et al. (2010)). In Section 7 we show that our model can be reinterpreted to the above framework, and this yields interesting insights about the value of information.

3 Model

Throughout the paper we focus on two-player games and we use index i to denote one of the two players and $-i$ to denote his opponent. An environment is a finite set of possible two-player games and a probability function according to which nature chooses the realized game. Formally:

Definition 1. An *environment* $((\Omega, p), A, u)$ is a tuple where:

- (Ω, p) is a finite probability space. That is, Ω is a finite set of states of nature, and p is a probability function over Ω. Without loss of generality, we assume that $p(\omega) > 0$ for each $\omega \in \Omega$.
- $A = A^1 \times A^2$, and each A^i is a finite set of pure actions for player i that are assumed to be available to him at all states of nature.[6]
- $u = \Omega \times A \to \mathbb{R}^2$ is the payoff function that determines the payoff for each player at each state of the world for each action profile.

We assume that both players evaluate random outcomes by using the von Neumann–Morgenstern expected utility extension of u, and, with a slight abuse of notation, we denote by u also this linear extension. Observe that each state of nature ω induces a normal-form game $(A, u(\omega))$ with a set of actions A and a payoff function u.

Remark 1. For simplicity, we define only for environments with two players. However, all the definitions and all the results can be extended to n-player environments in a straightforward way.

A *partition* over a set X is a set of non-empty subsets (*partition elements*) of X such that every element in X is in exactly one of these subsets. Let Π_Ω denote the set of all partitions over Ω. A (behavior) strategy of player i is a function $s_i : \Omega \to \Delta(A)$ that assigns a mixed action to each state of nature. Strategy s_i is π_i-measurable (where $\pi_i \in \Pi_\Omega$) if $s_i(\omega) = s_i(\omega')$ for each two states in the same partition element in π_i. Let S_{i,π_i} be the set of π_i-measurable (behavior) strategies of player i. Let $\bar{\pi} = \{\{\omega\} | \omega \in \Omega\}$ denote the *finest partition*, and let $\underline{\pi} = \{\Omega\}$ denote the *coarsest partition.* Given a partition $\pi_i \in \Pi_\Omega$ and a state $\omega \in \Omega$, let $\pi_i(\omega)$ be the partition element of π_i that includes state ω.

Each environment induces a two-player two-stage *meta-game*:

1. At stage 1, each player i chooses a partition π_i.
2. At stage 2, nature chooses the state $\omega \in \Omega$ according to p, and each player i observes the chosen partition profile (π_1, π_2) and his partition element $\pi_i(\omega)$, and chooses an action $a^i \in A^i$.

Remark 2.

[6] The assumption that a player has the same set of actions in all states is without loss of generality. If, originally, player i has different actions in different states, this can be modeled by having A_i large enough, relabeling actions in different states, and having several actions yielding the same outcome in some states.

1. We interpret the partition chosen by a player as a behavioral phenomenon, and not necessarily as a cognitive limitation. If two games belong to the same element of the partition it does not necessarily mean that the player is unaware of the fact that these are two different games. It does, however, say that the player is bounded to use the same strategy in both games. This could for instance arise from attributing irrelevance to the difference between these games or because the player is driven by some moral concern that leads him to use the same actions in both games. Note that this interpretation of the role of the partition allows us to assume that while a player bundles two games in the same element of the partition, he can still reasonably believe that his opponent will treat these two games differently.

2. We view the observability of the opponent's partition as a reduced-form model to a richer setup where players either (1) communicate about the strategic setup before playing it, and this reveals information about the opponent's partition, or (2) observe each other's past behavior or a trait that is correlated with the partition. The assumption of perfect observability is taken for tractability, and it is extended to partial observability in Section 7.2.

A meta-strategy is a behavior strategy in the meta-game.

Definition 2. A *(behavior) meta-strategy* for player i is a pair $\tau_i = (\mu_i, \sigma_i)$, where:

1. $\mu_i \in \Delta(\Pi_\Omega)$ is a distribution over the set of partitions.

2. $\sigma_i : \Pi_\Omega \times \Pi_\Omega \to S_{i,\pi_i}$ is a function that assigns a π_i-*measurable* (behavioral) strategy for each pair of partitions (π_i, π_{-i}).

Given a meta-strategy profile $\tau = (\tau_1, \tau_2)$ and a state $\omega \in \Omega$, let $x_\tau(\omega) \in \Delta(A)$ be the *interim outcome* of τ - the (possibly correlated) action profile that is induced in state ω given that both players follow the meta-strategy profile τ. Let $\bar{x}_\tau \in \Delta(A)$ be the *ex-ante outcome of* τ, i.e., $\forall a \in A$ $\bar{x}_\tau(a) = \sum_{\omega\in\Omega} p(\omega) \cdot x_\tau(\omega)(a)$. Let $C(\mu)$ denote the support of the distribution μ. With some abuse of notation, let π_i also denote the distribution with mass 1 on partition π_i.

A rule-rational equilibrium is a subgame-perfect equilibrium of the two-stage meta-game. Formally:

Definition 3. A meta-strategy profile $\tau^* = ((\mu_1^*, \sigma_1^*), (\mu_2^*, \sigma_2^*))$ is a *rule-rational equilibrium* if for each player i:

1. Each strategy is a best reply in each subgame: for each partition profile (π_1, π_2) and for each π_i-measurable strategy s_i':

$$\sum_{\omega\in\Omega} p(\omega) \cdot u_i\left(\sigma_i^*(\pi_1,\pi_2)(\omega), \sigma_{-i}^*(\pi_1,\pi_2)(\omega)\right) \geq \sum_{\omega\in\Omega} p(\omega) \cdot u_i\left(s_i'(\omega), \sigma_{-i}^*(\pi_1,\pi_2)(\omega)\right).$$

2. Each partition is a best reply: for each chosen partition $\pi_i^* \in C(\mu_i^*)$ and for each partition π_i':

$$\sum_{\pi_{-i}\in C\left(\mu_{-i}^*\right)} \mu_{-i}(\pi_{-i}) \cdot \sum_{\omega\in\Omega} p(\omega) \cdot u_i\left(\sigma_i^*(\pi_i^*,\pi_{-i})(\omega), \sigma_{-i}^*(\pi_1^*,\pi_{-i})(\omega)\right) \quad \geq$$
$$\sum_{\pi_{-i}^*\in C\left(\mu_{-i}^*\right)} \mu_{-i}(\pi_{-i}) \cdot \sum_{\omega\in\Omega} p(\omega) \cdot u_i\left(\sigma_i^*(\pi_i',\pi_{-i})(\omega), \sigma_{-i}^*(\pi_i',\pi_{-i})(\omega)\right).$$

Note that existence of rule-rational equilibria in any environment is immediately implied from the well-known existence result of subgame-perfect equilibria in finite games (Selten (1975)).

Remark 3.

1. We do not interpret the rule-rational equilibria to be the result of an explicit payoff maximization of "hyper-rational" agents who choose optimal partitions. Rather, we see the optimization in both stages as arising from an evolutionary process of social learning (similar to the literature of the "indirect evolutionary approach"; see, e.g., Dekel et al., 2007). Each agent in the population has a type that determines both his rules of thumb for bundling games together in a partition, and also his behavior in the second stage (given the partition). New agents usually adopt one of the existing types, and they are more likely to initiate the types of more successful agents. Every so often, a few agents ("mutants") may experiment with a new rule of thumb. It is well known that in the long run, the surviving types that are used in a stable state are optimal (i.e, Nash equilibria; see Nachbar, 1990). Recurrent entries of a few "mutants" and "trembles" of the incumbents imply that behavior also off-the-equilibrium path is optimal (i.e., subgame perfect).[7]

2. A rule-rational equilibrium differs from any standard solution concept of the incomplete information game induced by the environment (with a fixed information structure), as it allows players' plays to differ from best-replying in some states of nature. However, a rule-rational equilibrium can be interpreted as a solution concept of the incomplete information game in which players can choose (at the ex-ante stage) how much information to obtain (see Section 7.1 for further discussion).

The following definition divides the rule-rational equilibria into two disjoint subsets: (1) act-rational equilibria in which both players play a Nash equilibrium in all states, and (2) non-act-rational equilibria in which at least one of the players does not best reply in at least one state.

Definition 4. A rule-rational equilibrium $((\mu_1^*, \sigma_1^*), (\mu_2^*, \sigma_2^*))$ is *act-rational* if for each state $\omega \in \Omega$, and for each chosen partition profile $(\pi_1^*, \pi_2^*) \in C(\mu_1^*) \times C(\mu_2^*)$, the strategy profile $(\sigma_1^*(\pi_1^*, \pi_2^*)(\omega), \sigma_2^*(\pi_1^*, \pi_2^*)(\omega))$ is a Nash equilibrium in the game $(A, u(\omega))$. A rule-rational equilibrium is *non-act-rational* if it is not act-rational.

Remark 4.

1. Note that in an act-rational equilibrium a player is not required to play different actions in different games, or to use the finest partition. Moreover, the induced interim outcome $x_\tau(\omega)$ need not be a Nash equilibrium. However, it is immediate that the interim outcome in all states must be in the convex hull of the set of Nash equilibria (because the realized partition profile may be used as a random public signal to choose among the different Nash equilibria). Thus any act-rational equilibrium can be implemented with each player choosing his finest partition, provided that players have a rich enough source of "sunspots" to coordinate which Nash equilibrium they play.

2. In any non-act-rational equilibrium at least one of the players must not know his own payoff matrix in some states. That is, it must be that there exist two states ω, ω' and player $i \in \{1, 2\}$, such that with positive probability ω and ω' are in the same partition element of player i and $u_i(\omega) \neq u_i(\omega')$. Otherwise, rule rationality would imply act rationality.

[7]It is well known that any evolutionarily stable strategy (Maynard-Smith & Price (1973)) is a subgame-perfect equilibrium (see, e.g., Weibull, 1997). Evolutionary stability refines subgame perfection by also requiring stability against a small group of "mutants." We choose to use subgame perfection as the solution concept rather than evolutionary stability (or one of its adaptations to extensive-form games, such as Selten, 1983's notion of limit ESS), for the following reasons: (1) a subgame-perfect equilibrium is a simpler solution concept that is more familiar to readers and it suffices to yield most of the insights of the analysis, and (2) unlike evolutionarily stable strategies, subgame-perfect equilibria exist in all environments. We leave for future research the application of evolutionary stability in this setup.

4 Motivating Examples

This section includes four motivating examples. Example 2 shows an environment in which a coarse partition that pays attention only to the sum of payoffs yields a Pareto-dominant outcome in all states (relative to the set of act-rational payoffs). Example 3 demonstrates an environment that does not admit act-rational equilibria. Finally, in Section 4.3 we demonstrate that non-act-rational equilibria can fit experimentally observed behavior in the "Ultimatum game" (Example 4) and the "Chainstore paradox" (Example 5).

4.1 Non-act Rationality Pareto-dominates Act Rationality

Our first example demonstrates that non-act-rational equilibria may yield a strict Pareto improvement in all states relative to the best outcome that can be achieved by act-rational equilibria .

Example 2. Let $\Omega = \{\omega_1, \omega_2\}$, and let A and each ω_i be described as in Table 2.

Table 2: Payoffs of Prisoner's Dilemma with Two Defections

ω_1			
	C	D_1	D_2
C	9,9	0,10	2,8
D_1	10,0	1,1	1,1
D_2	8,2	1,1	1,1

ω_2			
	C	D_1	D_2
C	9,9	2,8	0,10
D_1	8,2	1,1	1,1
D_2	10,0	1,1	1,1

The interpretation of this environment is a Prisoner's Dilemma with two options to "testify against the partner" (D_1 and D_2), such that in each state of the world one of these actions is better for the defector. Observe that both games $(A, u(\omega_1))$ and $(A, u(\omega_2))$ admit a unique Nash equilibrium payoff of $(1, 1)$. This implies that $(1, 1)$ is the unique act-rational equilibrium payoff. This payoff can be implemented by the following act-rational equilibrium: both players choose the finest partition, and both play $0.5 \cdot D_1 + 0.5 \cdot D_2$ in both states (for any realized partition profile). One can see that the following meta-strategy profile is a rule-rational equilibrium for every $\frac{1}{10} < p(\omega_1) < \frac{9}{10}$:

- Coarse partitions for each player ($\pi_i = \underline{\pi}$).
- Playing (C, C) on the equilibrium path.
- Both players playing $0.5 \cdot D_1 + 0.5 \cdot D_2$ if any player deviated by choosing a different partition.

Observe that this rule-rational equilibrium induces a payoff of 9 for each player, which strictly Pareto-dominates the unique act-rational equilibrium payoff of 1 in both games. Further observe that the coarse partition in this example can be interpreted as if each player knows only the sum of the payoffs (which is the same in both states), but not how this sum is divided among the players when exactly one of the players plays C.

4.2 An Environment Without Act-rational equilibria

The following example demonstrates an environment that does not admit any act-rational equilibrium.

Example 3. Let $\Omega = \{\omega_1, \omega_2\}$, and let A and each $u(\omega_i)$ be described in Table 3.

Table 3: Environment with Act-Rational Equilibria

	ω_1	
	L	R
C	9,9	0,7
D_1	10,0	1,1
D_2	7,0	-2,1

	ω_2	
	L	R
C	9,9	0,7
D_1	7,0	-2,1
D_2	10,0	1,1

We now show that for every $\frac{1}{3} < p(\omega_1) < \frac{2}{3}$, the game admits no act-rational equilibrium. This implies that this is what the players must play in any act-rational equilibrium, and that each player obtains a payoff of 1. However, if player 1 deviates to the coarsest partition $\underline{\pi}$, then for both possible partitions of player 2, the induced Bayesian subgame admits a unique Nash equilibrium in which the players play (C, L) and this yields both players a payoff of 9. This shows that the environment does not admit any act-rational equilibrium. An example for a non-act-rational equilibrium in this environment is as follows: let $\pi_1^* = \underline{\pi}$, $\pi_2^* = \overline{\pi}$, and let σ^* be equal to (C, L) on-the-equilibrium path or if player 2 deviated to $\underline{\pi}$; otherwise (i.e., if player 1 deviated to $\overline{\pi}$) let σ^* be equal to (D_1, R) in ω_1 and (D_2, R) in ω_2 . Observe that each game ω_i admits a unique Nash equilibrium: (D_1, R) in ω_1 and (D_2, R) in ω_2.

4.3 Environments with Extensive-form Games

To simplify the presentation of the general model in Section 3 we defined an environment consisting of normal-form games. The following two examples differ somewhat in that in each state of nature players interact in sequential games. Accordingly, with some abuse of notation, in this subsection we have extensive-form games in each state of nature, and we require a rule-rational equilibrium to be a subgame-perfect equilibrium (given the extensive-form representation of the games).

4.3.1 Ultimatum Game

The following example describes a situation in which players may either play an Ultimatum game, or a bargaining game in which the responder can make a counteroffer.

Example 4. Consider an environment that includes two states $\Omega = \{\omega_1, \omega_2\}$ with equal probability. In both states the two risk-neutral[8] players have to agree how to divide 1 dollar between them. In ω_1 they play the Ultimatum game: player 1 offers a division, which player 2 either accepts or rejects; in the case of rejection both players get nothing. In ω_2 they play a two-stage sequential bargaining game: at stage 1 player 1 makes an offer; then player 2 either accepts the offer or makes a counteroffer; in the case of a counteroffer, player 1 may either accept or reject it; in the case of rejection both players get nothing.[9]

One can see that the game admits two pure rule-rational equilibria:

1. Act-rational equilibrium - Both players choose the finest partition. Player 1 offers 0 to the responder in ω_1 (Ultimatum) and 1 in ω_2. Player 2 accepts these offers (and rejects any lower offer off the equilibrium path).

2. Non-act-rational equilibrium - Player 2 chooses the coarsest partition. Player 1 may either choose the finest or the coarsest partition. Player 1 offers the responder 0.5 in both games. Player 2

[8] Results are qualitatively the same for risk-averse players and for having a discount factor.

[9] As stated above, the players have different sets of actions in each state. However, this can also be modeled with the same set of actions in each state. Specifically, one may assume that in both states the players have the richer set of actions of state ω_2, and that in state ω_1 the counteroffer of player 2 leads to a payoff of 0 for each player regardless of player 1's final action (accepting or rejecting).

accepts this offer (and rejects any lower offer off the equilibrium path).

Observe that both equilibria yield both players the same ex-ante payoff (0.5 for each player). The second equilibrium fits the stylized fact that people tend to reject offers that are substantially lower than half the pie when playing the ultimatum game.

4.3.2 Chainstore Paradox

The following example shows that a non-act-rational equilibrium, in which the monopolist chooses not to pay attention to the exact number of potential entrants, implies a reputation-like behavior that fits the experimentally observed behavior, which cannot be supported by act-rational equilibria. In this example we extend the model by allowing the set of states to be countable (and not finite as in our general model).

Example 5. Consider an environment with a countable number of states: $\Omega = \mathbb{N}$. Let p be any distribution satisfying: (1) full support - $p(\omega) > 0 \ \forall \omega \in \mathbb{N}$, and (2) "continuation" probability that is high enough - $\forall k \in \mathbb{N}$, $p(\omega = k | \omega \geq k) < 2/3$. In particular these requirements are satisfied by any geometric distribution with high enough expectation. The state ω describes the number of potential competitors that the *monopolist* (player 1) faces. The game lasts ω rounds. To simplify notation and fit our two-player setup, we represent all potential competitors by a single "impatient" player 2, the *competitor*. At each round, player 2 chooses In or Out. Out yields a stage payoff of 5 to the monopolist and 1 to the competitor. If he chose In, then the monopolist chooses Fight or Adapt. Fight yields stage payoff of 0 to both players, and Adapt yields 2 to both players. The total game payoff of the monopolist is the discounted sum of stage payoffs for some high enough discount factor $0 << \lambda \leq 1$. The total payoff of the competitor is a discounted sum of the payoffs with a low discount factor $0 \leq \lambda << 1$. In the classical setup in which the monopolist knows the state of nature, the game admits a unique subgame-perfect equilibrium: the competitor always plays In, and the monopolist always plays Adapt. In our setup, however, there is an additional non-act-rational equilibrium that yields a better payoff to the monopolist. In this equilibrium the monopolist chooses the coarsest partition, and plays Fight following any entry. The competitor chooses the finest partition and plays Out, unless he observes an out-of-equilibrium history in which the monopolist played Adapt.

5 Results for Specific Families of Environments

This section characterizes rule-rational equilibria in families of environments that satisfy certain properties. In Section 5.1 we show that rule-rational equilibria are closely related to correlated equilibria in environments that include strategically equivalent games. In Section 5.2 we show that environments that only include constant-sum games essentially admit only act-rational equilibria (but this is not true for the larger family of strategically zero-sum games). Section 5.3 shows that environments with common interests admit a Pareto-dominant act-rational equilibrium. In Section 5.4 we characterize rule-rational equilibria in environments in which a specific action profile is a Nash equilibrium in all games. Finally, in Section 5.5 we study environments that assign high enough probability to a game with an undominated, strict equilibrium.

5.1 Redundant Environments

Recall (Moulin & Vial, 1978), that two games with the same set of actions, (A, u) and $(A, \tilde{u})$, are *strategically equivalent* if for each player $i \in \{1, 2\}$, for each mixed action $x_i \in \Delta(A_i)$, and for each two mixed actions $x_{-i}, x'_{-i} \in \Delta(A_{-i})$, $u_i(x_i, x_{-i}) \geq u_i(x_i, x'_{-i}) \Leftrightarrow \tilde{u}_i(x_i, x_{-i}) \geq \tilde{u}_i(x_i, x'_{-i})$. That is, each player has the same preference over his set of mixed actions given any mixed action of his opponent. It is immediate to see that strategically equivalent games admit the same sets of Nash equilibria and correlated equilibria. An environment is *redundant* if it includes only strategically equivalent games. That is, for each $\omega, \omega' \in \Omega$, the games $(A, u(\omega))$ and $(A, u(\omega'))$ are strategically equivalent.

Recall that a correlated equilibrium is a Nash equilibrium of an extended game that includes a pre-play stage in which each player observes a private signal (and the signals may be correlated). Formally:

Definition 5. (Aumann, 1974, as reformulated in Osborne & Rubinstein, 1994) A correlated equilibrium of a two-player normal-form game $G = (A, u)$ is a tuple $\zeta = \left(\left(\tilde{\Omega}, \tilde{p}\right), (\tilde{\pi}_1, \tilde{\pi}_2), (\tilde{\eta}_1, \tilde{\eta}_2)\right)$ that consists of:

- A finite probability space $\left(\tilde{\Omega}, \tilde{p}\right)$,
- A partition profile $(\tilde{\pi}_1, \tilde{\pi}_2) \in \Pi_{\tilde{\Omega}} \times \Pi_{\tilde{\Omega}}$,
- A strategy profile $(\tilde{\eta}_1, \tilde{\eta}_2)$, where each $\tilde{\eta}_i : \Omega \to A_i$ is a $\tilde{\pi}_i$-measurable strategy,

such that for every player i and every π_i-measurable strategy $\eta'_i : \Omega \to A_i$, we have

$$\sum_{\omega \in \tilde{\Omega}} \tilde{p}(\omega) \cdot u_i(\tilde{\eta}_i(\omega), \tilde{\eta}_{-i}(\omega)) \geq \sum_{\omega \in \tilde{\Omega}} \tilde{p}(\omega) \cdot u_i(\eta'_i(\omega), \tilde{\eta}_{-i}(\omega)).$$

The following example demonstrates that a redundant environment that includes only copies of the same game, may admit a rule-rational equilibrium with a payoff outside the convex hull of the set of Nash equilibrium payoffs.

Example 6. Consider the following redundant environment: $\Omega = \{\omega_1, \omega_2, \omega_3\}$, p is uniform, and (A, u) is the same "Chicken" (or "Hawk-Dove") game in all states (see Table 4). Consider the following rule-

Table 4: Payoffs of the "Chicken" Game in All States

	C	D
C	3,3	1,4
D	4,1	0,0

rational equilibrium: player 1 chooses partition $\pi_1^* = \{\{\omega_1, \omega_2,\} \omega_3\}$ and player 2 chooses partition $\pi_2^* = \{\omega_1, \{\omega_2, \omega_3\}\}$. On the equilibrium path each player plays C when he does not know the exact state of nature, and plays D otherwise. Off the equilibrium path, if a player chooses any other partition, the deviator plays C and his opponent plays D.[10] This rule-rational equilibrium (which "imitates" a correlated equilibrium of the "Chicken" game) yields each player a payoff of 8/3 (and this is outside the convex hull of the set of Nash equilibria).

The following proposition shows that redundant environments can only induce behavior that is consistent with correlated equilibria. Formally:

Proposition 1. *Let $((\Omega, p), A, u)$ be a redundant environment. Then:*

1. *Any Nash equilibrium can be induced by an act-rational equilibrium.*

[10]For brevity, throughout the proofs in this paper we specify only the strategies after a single player deviates. The play following simultaneous deviations of both players at the partition stage can be determined arbitrarily as any Nash equilibrium of the induced subgame.

2. *The induced aggregate action profile of every rule-rational equilibrium is always a correlated equilibrium (of every game in the environment).*

Proof. □

1. Let $x \in \Delta(A_1) \times \Delta(A_2)$ be a Nash equilibrium of the games $(A, u(\omega))_{\omega \in \Omega}$. It is immediate that any arbitrary choice of partitions, followed by playing x in all states (regardless of the chosen partitions), yields an act-rational equilibrium.

2. Let $\tau^* = ((\mu_1^*, \sigma_1^*), (\mu_2^*, \sigma_2^*))$ be a rule-rational equilibrium. Intuitively, a mediator can "mimic" the choice of the state of nature and of the partition that each player uses, and inform each player about his partition and about the partition element that includes the state. Formally, we construct a correlated equilibrium $\zeta_{\tau^*} = \left(\left(\tilde{\Omega}, \tilde{p}\right), (\tilde{\pi}_1, \tilde{\pi}_2), (\tilde{\eta}_1, \tilde{\eta}_2)\right)$ (of every game in the environment) as follows. Let $\tilde{\Omega} = \Omega \times \Pi_\Omega \times \Pi_\Omega$, and let $\tilde{p}(\omega, \pi_1, \pi_2) = p(\omega) \cdot \mu_1^*(\pi_1) \cdot \mu_2^*(\pi_2)$. For each player i, let $\tilde{\pi}_i((\omega_0, \pi_1, \pi_2)) = \{(\omega, \pi_1, \pi_2) | \omega \in \pi_i(\omega_0)\}$, and let $\tilde{\eta}_i(\omega_0) = \sigma_i^*(\pi_1, \pi_2)(\omega_0)$. The fact that τ^* is a rule-rational equilibrium immediately implies that ζ_{τ^*} is a correlated equilibrium.

The following proposition shows that the converse is also true: any correlated equilibrium (which is not dominated by the set of Nash equilibria) can be induced by a redundant environment that includes several copies of the same game.

Definition 6. A correlated equilibrium is undominated if for each player i there exists a Nash equilibrium that yields player i a weakly worse payoff.

Proposition 2. *For every undominated correlated equilibrium ζ in game (A, u_0), there exists a redundant environment $((\Omega, p), A, u)$ and a rule-rational equilibrium that induces ζ as the aggregate outcome.*

Proof. Let $\zeta = \left(\left(\tilde{\Omega}, \tilde{p}\right), (\tilde{\pi}_1, \tilde{\pi}_2), (\eta_1, \eta_2)\right)$ be a correlated equilibrium of the game (A, u_0). Consider the following redundant environment $((\Omega, p), A, u)$: $(\Omega, p) = \left(\tilde{\Omega}, \tilde{p}\right)$ and $u(\omega) = u_0$ in all states. Consider the following strategy profile $\tau_\zeta = ((\mu_1, \sigma_1), (\mu_2, \sigma_2))$: each μ_i assigns mass 1 to $\tilde{\pi}_i$, and $\sigma_i = \eta_i$ if the realized partition profile is $(\tilde{\pi}_1, \tilde{\pi}_2)$. If player i deviates and plays a different partition, both players play in all states the Nash equilibrium that yields player i a worse payoff. It is immediate to see that τ_ζ is a rule-rational equilibrium.

Finally, the following proposition shows that the implementation of correlated equilibria does not depend on using "tailored" probability distributions. Specifically, it shows that any undominated correlated equilibrium can be approximately induced by any large enough "atomless" environment (an environment in which each state has a small probability). Formally: □

Definition 7. Given $\epsilon \geq 0$, a meta-strategy profile $\tau^* = ((\mu_1^*, \sigma_1^*), (\mu_2^*, \sigma_2^*))$ is a *rule-rational ϵ-equilibrium* if for each player i:

1. Each strategy is an ϵ-best reply in each sub-game: for each partition-profile (π_1, π_2) and for each π_i-measurable strategy s_i':

$$\sum_{\omega \in \Omega} p(\omega) \cdot u_i\left(\sigma_i^*(\pi_1, \pi_2)(\omega), \sigma_{-i}^*(\pi_1, \pi_2)(\omega)\right) \geq \sum_{\omega \in \Omega} p(\omega) \cdot u_i\left(s_i'(\omega), \sigma_{-i}^*(\pi_1, \pi_2)(\omega)\right) - \epsilon.$$

2. Each partition is an ϵ-best reply: for each chosen partition $\pi_i^* \in C(\mu_i^*)$ and for each partition π_i':

$$\sum_{\pi_{-i} \in C\left(\mu_{-i}^*\right)} \mu_{-i}(\pi_{-i}) \cdot \sum_{\omega \in \Omega} p(\omega) \cdot u_i\left(\sigma_i^*(\pi_i^*, \pi_{-i})(\omega), \sigma_{-i}^*(\pi_1^*, \pi_{-i})(\omega)\right) \geq$$
$$\sum_{\pi_{-i}^* \in C\left(\mu_{-i}^*\right)} \mu_{-i}(\pi_{-i}) \cdot \sum_{\omega \in \Omega} p(\omega) \cdot u_i\left(\sigma_i^*(\pi_i', \pi_{-i})(\omega), \sigma_{-i}^*(\pi_i', \pi_{-i})(\omega)\right) - \epsilon.$$

Definition 8. Given a normal-form game (A, u) and $\epsilon \geq 0$, two correlated action profiles $x, y \in \Delta(A)$ are ϵ-close if $\max_{a \in A}(|x(a) - y(a)|) \leq \epsilon$.

Definition 9. Given a normal-form game (A, u_0), let $\bar{u}_0 = max_{a,a' \in A, i \in \{1,2\}} |u_i(a) - u_i(a')|$ be the maximal difference between any two values in the payoff matrix.

Proposition 3. *Let (A, u_0) be a game, let ζ be an undominated correlated equilibrium in (A, u_0), and let $\epsilon > 0$. Let $((\Omega, p), A, u_0)$ be a redundant environment that satisfies $p(\omega) < \epsilon/\bar{u}_0$ $\forall \omega \in \Omega$. Then, the environment $((\Omega, p), A, u_0)$ admits a rule-rational ϵ-equilibrium that induces an ex-ante outcome, which is ϵ-close to ζ.*

Proof. Let $x_\zeta \in \Delta(A)$ be the outcome of the correlated equilibrium ζ. It is immediate that there exists a function $g : \Omega \to A$ that identifies each set of states with an action profile in A, such that for each action profile $a \in A$, $|p(\{\omega | g(\omega) = a\}) - x_\zeta(a)| < \epsilon/\bar{u}_0$. Let $g_i(\omega) : \Omega \to A_i$ be the function that identifies each state with the induced action of player i (i.e, $g(\omega) = (g_1(\omega), g_2(\omega))$. For each player i let partition π_i be defined as follows: $\pi_i(\omega_0) = \{\omega | g_i(\omega) = g_i(\omega_0)\}$. Consider the following meta-strategy profile $\tau_{\zeta,\epsilon} = ((\pi_1, \sigma_1), (\pi_2, \sigma_2))$: on the equilibrium path $\sigma_i(\pi_1, \pi_2)(\omega) = g_i(\omega)$, and if player i deviates and chooses a different partition, then both agents play in all states the Nash equilibrium that yields player i a worse payoff. It is immediate to see that $\tau_{\zeta,\epsilon}$ is a rule-rational ϵ-equilibrium that yields an outcome that is $\epsilon/\bar{u}_0$ close to ζ. □

5.2 Constant-sum and Strategically Zero-sum Environments

Environment $((\Omega, p), A, u)$ is constant-sum if for each state $\omega \in \Omega$, the game $(A, u(\omega))$ is a constant-sum game. For each state $\omega \in \Omega$, let $v(\omega)$ be the interim value of the game $(A, u(\omega))$, and let $v(\Omega) = \sum p(\omega) \cdot v(\omega)$ be the *ex-ante* value of the environment. The following example shows a zero-sum environment that admits a non-act-rational equilibrium.

Example 7. Let $\Omega = \{\omega_1, \omega_2\}$, $p(\omega_1) = 0.5$, and let A and each $u(\omega)$ be defined according to Table 5. Consider the following meta-strategy profile. Both players choose the coarsest partition ($\forall i$ $\pi_i^* = \underline{\pi}$)

Table 5: Payoffs of a Zero-sum Environment

ω_1		
	L	R
T	1,-1	0,0
D	0,0	0,0

ω_2		
	L	R
T	-1,1	0,0
D	0,0	0,0

and they both play (T, L) on the equilibrium path, and (R, D) off the equilibrium path (if any player deviated and chose a different partition). It is immediate to see that this profile is a non-act-rational equilibrium. Note that this equilibrium yields the players in each state ω_i a payoff different from the interim value $v(\omega_i) = 0$, while the ex-ante value, $v(\Omega) = 0$, is unchanged.

Our next result shows that rule rationality is very limited in constant-sum environments. Specifically, any rule-rational equilibrium induces the same *ex-ante* payoff as the act-rational equilibria and it is not robust, in the sense that stability is destroyed if there is an arbitrarily small probability that the opponent may overlook a deviation at the partition stage. Formally, our notion of robustness assumes that when a player deviates while choosing a partition, then: (1) with very high probability the deviation is observed by the opponent and the players follow the off-equilibrium-path strategies, and (2) with very low probability the deviation is unobserved by the opponent, and the deviating player (who is aware of the fact that his deviation was unobserved) can play any strategy that is consistent with his partition.

Definition 10. *Rule-rational equilibrium* $\tau^* = ((\mu_1^*, \sigma_1^*), (\mu_2^*, \sigma_2^*))$ *is robust if there exists* $\bar{\epsilon} > 0$ *such that for each* $\epsilon < \bar{\epsilon}$, for each player i, for each chosen partition $\pi_i^* \in C(\mu_i^*)$, and for each partition π_i':

$$\sum_{\pi_{-i}^* \in C\left(\mu_{-i}^*\right)} \mu_{-i}(\pi_{-i}) \cdot \sum_{\omega \in \Omega} p(\omega) \cdot u_i\left(\sigma_i^*(\pi_1^*, \pi_2^*)(\omega), \sigma_{-i}^*(\pi_1^*, \pi_2^*)(\omega)\right) \geq$$

$$\sum_{\pi_{-i}^* \in C\left(\mu_{-i}^*\right)} \mu_{-i}(\pi_{-i}) \cdot \left((1-\epsilon) \cdot \sum_{\omega \in \Omega} p(\omega) \cdot u_i\left(\sigma_i^*\left(\pi_i', \pi_{-i}^*\right)(\omega), \sigma_{-i}^*\left(\pi_i', \pi_{-i}^*\right)(\omega)\right)\right.$$

$$\left. + \epsilon \cdot \sum_{\omega \in \Omega} p(\omega) \cdot u_i\left(\max_{\sigma_i \in S_{i,\pi'_i}} \sigma_i(\omega), \sigma_{-i}^*(\pi_1^*, \pi_2^*)(\omega)\right)\right).$$

Proposition 4. *Let* $((\Omega, p), A, u)$ *be a constant-sum environment. Then:*

1. *The environment admits an act-rational equilibrium.*
2. *All rule-rational equilibria yield the same ex-ante payoff,* $v(\Omega)$.
3. *Any non-act-rational equilibrium is not robust.*

Proof.

1. Consider the meta-strategy in which each player chooses the finest partition and then plays a minimax strategy in each game (regardless of the partition his opponent chose). It it immediate that this meta-strategy profile is an act-rational equilibrium.
2. This is immediately implied by the fact each player may deviate by choosing the finest partition, and playing a minimax strategy in each game, which yields an *ex-ante* payoff of at least $v(\Omega)$.
3. For every $\epsilon > 0$, a player can strictly earn by deviating from a rule-rational equilibrium into choosing the finest partition, and playing a minimax strategy in each game. The deviator will earn the same *ex-ante* payoff if the deviation is observed, and will earn a strictly higher payoff otherwise (as there is a state in which with positive probability the players do not play a Nash equilibrium, and one of the player can earn by best-replying to its opponent's play).

□

Recall (Moulin & Vial, 1978) that a game is strategically zero-sum if it is strategically equivalent to a zero-sum game. An environment $((\Omega, p), A, u)$ is strategically zero-sum if for each state $\omega \in \Omega$, the game $(A, u(\omega))$ is strategically zero-sum. Note that each constant-sum environment is strategically zero-sum, but the converse is not necessarily true. The following two examples demonstrate that all parts of Proposition 4 may fail in strategically zero-sum environments. Example 8 demonstrates a strategically zero-sum environment with a robust non-act-rational equilibrium that ex-ante strictly Pareto-dominates the unique act-rational equilibrium payoff.

Example 8. Let $\Omega = \{\omega_1, \omega_2\}$, $p(\omega_1) = 0.5$, and let A and each $u(\omega)$ be defined as in Table 6. First,

Table 6: Payoffs of a Strategically Zero-sum Environment

ω_1		
	L	R
T	2,-1	0,0
D	0,0	0,0

ω_2		
	L	R
T	-1,2	0,0
D	0,0	0,0

note that each game is strategically zero-sum. Specifically, the first (second) game $A, u(\omega_1)$ $(A, u(\omega_2))$ is derived from the first (second) zero-sum game of Example 7 by multiplying the payoff matrix of player 1 (player 2) by two. Next, note that each game admits a unique Nash equilibrium payoff of $(0,0)$. This implies that any act-rational equilibrium admits an interim payoff $(0,0)$ in both states. Finally, consider the following meta-strategy profile. Both players choose the coarsest partition ($\forall i\ \pi_i^* = \underline{\pi}$) and they both play (T, L) on the equilibrium path, and (R, D) off the equilibrium path (if any player has deviated and chosen a different partition). It is immediate that this profile is a non-act-rational equilibrium, and that it yields each player an ex-ante payoff of $(0.5, 0.5)$, which Pareto-dominates the payoff of each act-rational equilibrium.

The following example demonstrates a strategically zero-sum environment that does not admit any act-rational equilibrium.

Example 9. Let $\Omega = \{\omega_1, \omega_2\}$, $p(\omega_1) = 0.5$, and let A and each $u(\omega)$ be defined as in Table 7. First,

Table 7: Payoffs of a Strategically-Zero-Sum Environment

ω_1		
	L	R
T	0,0	0,0
D	3,-3	4,-4

ω_2		
	L	R
T	11,-1	-1,1
D	9,1	1,-1

note that each game is strategically zero-sum: $(A, u(\omega_1))$ is a zero-sum game, and $(A, u(\omega_2))$ is derived from a zero-sum Matching Pennies game by adding 10 to the first column of the matrix payoff of player 1 (Moulin & Vial, 1978 show that this implies that $(A, u(\omega_2))$ is strategically zero-sum). Next, note that each game admits a unique Nash equilibrium: (L, D) in ω_1 and $((0.5, 0.5), (0.5, 0.5))$ in ω_2. These action profiles yield player 1 an ex-ante payoff of $4 = 0.5 * 3 + 0.5 * 5$. Assume to the contrary that these action profiles were the interim outcome of an act-rational equilibrium. Note that player 1 can deviate by choosing the coarsest partition. Given this deviation, playing D is a strictly dominant action for player 1, and this implies that playing(D, L) is the unique Nash equilibrium of the induced game, and that it yields player 1 an ex-ante payoff of 6. This contradicts the assumption that an act-rational equilibrium exists.

5.3 Environments with Common Interests

A game (A, u_0) has *common interests* (Aumann & Sorin, 1989) if there is an action profile $\tilde{a} \in A$ that Pareto-dominates all other action profiles: $\forall i \in \{1,2\}, a \in A, u_i(\tilde{a}) > u_i(a)$. An environment $((\Omega, p), A, u)$ has common interests if it includes only games with common interests. The following example shows that rule-rational equilibria may exist in environments with common interests.

Example 10. Let $\Omega = \{\omega_1, \omega_2\}$, and let A and each $u(\omega_i)$ be defined as in Table 8. One can see that

Table 8: Payoffs of an Environment with Common Interests

ω_1				ω_2		
	L	R			L	R
T	5,5	0,0		T	5,5	0,0
M	6,6	0,0		M	0,0	0,0
D	0,0	1,1		D	0,0	1,1

the following meta-strategy profile is a non-act-rational equilibrium for every $p(\omega_1) < 5/6$: the players choose the coarsest partition ($\pi_i = \underline{\pi}$) and play (T, L) on the equilibrium path, and (D, R) if any player deviated by choosing a different partition.

The following proposition shows that rule rationality is very limited in environments with common interests, in the sense that any rule-rational equilibrium is strictly Pareto-dominated by an act-rational equilibrium.[11] The proposition is related to an existing result of Bassan et al. (2003), which shows that information has a (Pareto) positive value if each game has common interests.

Proposition 5. *Let $((\Omega, p), A, u)$ be an environment with common interests. Then there exists an act-rational equilibrium that strictly Pareto-dominates all non-act-rational equilibria.*

Proof. It is immediate that the meta-strategy profile in which players choose the finest partition and then play the Pareto-optimal action profile in every game is an act-rational equilibrium. Consider any non-act-rational equilibrium. This equilibrium induces in some $\omega_0 \in \Omega$ a profile that is not a Nash equilibrium. The fact that each game has common interests implies that the interim payoff of the non-act-rational equilibrium is weakly dominated by the interim payoff of the above act-rational equilibrium, and that it is strictly dominated in ω_0. □

5.4 Uniform-profile Environments

The following simple claim deals with environments with a uniform Nash equilibrium, i.e., with environments in which the same mixed action profile is a Nash equilibrium in all states. It shows that such environments admit an act-rational equilibrium in which the above action profile is played in all states. Formally:

Claim 1. Let $((\Omega, p), A, u)$ be an environment, and let $x \in \Delta(A)$ be a mixed action profile that is a Nash equilibrium in all states. The the environment admits an act-rational equilibrium in which the players choose the coarsest partitions and play x in all states.

Proof. It is immediate that the following meta-strategy profile is a rule-rational equilibrium: both players choose the coarsest partition $\underline{\pi}$, and play mixed action x regardless of the realized partition profile. □

The following proposition shows that if an environment admits a uniform dominant action for one of the players, then the environment admits *only* act-rational equilibria. Formally:

Proposition 6. *Let $((\Omega, p), A, u)$ be an environment with a player $i \in \{1, 2\}$ and an action $a^i \in A^i$ that is strictly dominant in all states. Then the environment admits only act-rational equilibria.*

[11] In addition, one can show that all the non-act-rational equilibria in common-interest environments do not satisfy the forward-induction refinement (Govindan & Wilson (2008)), while the Pareto-dominant act-rational equilibrium satisfies it.

Proof. Let $a_i \in A_i$ be the strictly dominant action for player i in all games. It is immediate that the following meta-strategy profile is an act-rational equilibrium: both players choose the finest partition, player i plays a_i in each state (also off the equilibrium path), and player $-i$ best-replies to action a_i in each game. Any strategy profile that induces a different interim outcome than playing a_i by player i and best-replying to a_i in each state by player $-i$ is not a rule-rational equilibrium because: (1) if player i plays with positive probability a different action in some state, then he can strictly earn more by deviating and playing a_i in all states; and (2) if player i always plays a_i but his opponent does not best-reply to a_i in some state, then player $-i$ can earn strictly more by deviating to the finest partition and best-replying to a_i in all states. □

5.5 Environments with an Undominated Strict Equilibrium

In this section we consider an environment that includes a game g^* with a strict equilibrium a that dominates other strict equilibria. The first result shows that if the other games in the environment have payoff matrices that are close enough to g^*, and if the probability of g^* is large enough, then the environment admits a rule-rational equilibrium in which players bundle all the similar games together by choosing the coarsest partition and then playing a in all states. Note that if there is a state in which a is not a Nash equilibrium, then this is a non-act-rational equilibrium. Formally:

Definition 11. Let $\delta > 0$. A pure action profile a in a game (A, u) is a δ-*strict Nash equilibrium* if for each player i and each action $b^i \in A^i$, $u^i(a) \geq u^i\left(b^i, a^{-i}\right) + \delta$.

Proposition 7. *Let $\epsilon, \delta > 0$. Let (A, u_0) be a game with δ-strict equilibria (a, b_1, b_2) such that $\forall i \in \{1, 2\}$, $u_i(a) > u_i(b_i) + \delta$. Let $((\Omega, p), A, u)$ be an environment such that: (1) it includes game (A, u_0): $\exists \omega_0 \in \Omega$ such that $u(\omega_0) = u_0$; (2) the payoff matrices in all states are ϵ-close to u_o: $\forall \omega \in \Omega$, $a \in A\, i \in \{1, 2\}$: $\left|u^i(\omega)(a) - u_0^i(a)\right| \leq \epsilon$; and (3) $p(\omega_0) \geq \frac{\epsilon - \delta}{\epsilon}$. Then the environment admits a rule-rational equilibrium in which players choose the coarsest partition and play a in all states.*

Proof. Consider the following meta-strategy profile. The players choose the coarsest partition and play a. If player i deviates and chooses a different partition, then the players play profile b_i in the partition element that includes ω_0, and they play an arbitrary equilibrium in the subgame that is induced by the choice of partitions and by fixing the play of profile b_i in ω_0. Note that any unilateral deviation in the partition element that includes ω_0 yields the deviator a loss of at least δ utility points in ω_0 and a profit of at most $\epsilon - \delta$ in any other state. Thus, this deviation yields an expected profit of at most

$$\left(1 - \frac{\epsilon - \delta}{\epsilon}\right) \cdot (\epsilon - \delta) - \frac{\epsilon - \delta}{\delta} \cdot \delta = \frac{\delta}{\epsilon} \cdot \epsilon - \delta = 0.$$

This implies that the above strategy is an equilibrium of the subgame that is induced by the chosen partitions also without fixing profile b_i in ω_0. Observe that a player who deviates at the partition stage loses at least δ in ω_0 and gains at most $\epsilon - \delta$ in any other state; thus the maximal profit from a deviation from the meta-strategy profile is

$$\left(1 - \frac{\epsilon - \delta}{\epsilon}\right) \cdot (\epsilon - \delta) - \frac{\epsilon - \delta}{\epsilon} \cdot \delta = 0.$$

This implies that the above meta-strategy profile is a rule-rational equilibrium. □

The following corollary shows that if g^* (a game with a strict equilibrium a that dominates other strict equilibria) has high enough probability, then the environment admits a rule-rational equilibrium in which players bundle all the other "unlikely" games with the prominent game g^*, and play a in all games. Formally:

Corollary 1. *Let (A, u_0) be a game with strict equilibria (a, b_1, b_2) such that $\forall i \in \{1,2\}$, $u_i(a) > u_i(b_i)$. Then there exists $0 < p_0 < 1$ such that each environment $((\Omega, p), A, u)$ that satisfies $\omega_0 \in \Omega$, $u(\omega_0) = u_0$, and $p(\omega_0) \geq p_0$, admits a rule-rational equilibrium in which players choose the coarsest partition and play a in all states.*

Proof. Let $\delta > 0$ be small enough such that: (1) the profiles (a, b_1, b_2) are δ-strict equilibria in game (A, u_0), (2) for each player i, $u_i(a) - u_i(b_i) > \delta$. Let $\epsilon > 0$ be large enough such that all games are ϵ-close to g^*: $\forall \omega \in \Omega$, $a^1 \in A^1$, $a^2 \in A^2$ $i \in \{1,2\}$ $\left|u^i\left(a^1, a^2\right) - u_0^i\left(a^1, a^2\right)\right| \leq \epsilon$. Let $p_0 = \frac{\epsilon - \delta}{\epsilon}$. Then proposition 7 implies that there is a rule-rational equilibrium in which the players choose the coarsest partition, and play a in all states. □

6 Stackelberg Stability and Act Rationality

This section shows that the existence of act-rational equilibria is closely related to the question of whether the Nash equilibria of the games in the environment are stable against the possibility that one of the players becomes a "Stackelberg" leader and commits in advance to a specific strategy. At the end of Section 2 we discussed the implications of the results of this section for the study of the value of information.

Given mixed action $x^i \in \Delta\left(A^i\right)$ of player i in a normal-form game $G = (A, u)$, let $BR\left(x^i\right) \subseteq A^{-i}$ be the set of (pure) best replies of player $-i$ to x_i. Let $v^i\left(x^i\right) = \min_{x^{-i} \in BR(x^i)} u_i\left(x^i, x^{-i}\right)$ be the payoff that player i can guarantee himself when playing x^i before his opponent plays (under the assumption that his opponent best-replies to the player's commitment). Let $\bar{v}^i(G) = \max_{x^i \in \Delta(A^i)} v^i\left(x^i\right)$ be the payoff that player i can guarantee when playing as a Stackelberg leader (called a mixed-action Stackelberg payoff in Mailath & Samuelson (2006)). Let $NE(G)$ be the set of Nash equilibrium payoff vectors in G.

We say that a Nash equilibrium is Stackelberg-stable if no player can guarantee a higher payoff by being a Stackelberg leader. Formally:

Definition 12. A Nash equilibrium payoff vector $u \in NE(G)$ is *Stackelberg-stable* if for each player i, $\bar{v}_i(G) \leq u_i$.

A game is *Stackelberg-stable* if it admits at least one Stackelberg-stable equilibrium, and it is *Stackelberg-unstable* otherwise. The following fact shows a few simple families of Stackelberg-stable games. In particular, both constant-sum games and optimal-profile games have this property.

Fact 1. *Each of the following games are Stackelberg-stable.*

1. *Constant-sum games (all equilibria are stable).*
2. *Optimal-profile games (the Pareto-dominant equilibrium is stable).*
3. *Games in which both players have a strictly dominant action, e.g., Prisoner's Dilemma (the unique equilibrium is stable).*

The following propositions show the close relation between act-rational equilibria and Stackelberg-stable games. The first proposition shows that if all games are Stackelberg-stable, then an act-rational equilibrium exists.

Proposition 8. *If every game in the environment is Stackelberg-stable, then the environment admits an act-rational equilibrium.*

Proof. Consider the following meta-strategy profile. Each player chooses the finest partition, and the players play on the equilibrium path a Stackelberg-stable equilibrium in each game. If any player deviates to a different partition, then the players play a Nash equilibrium of the induced subgame, such that the deviating player earns weakly less than its payoff on the equilibrium path. The induced subgame admits such a Nash equilibrium because otherwise the deviating player would have a strictly profitable Stackelberg commitment in at least one game in the environment, and this would contradict our assumption that players play Stackelberg-stable equilibria in all states. The above argument implies that this meta-strategy profile is an act-rational equilibrium. □

The following proposition shows that a single Stackelberg-unstable equilibrium is enough to prevent the existence of an act-rational equilibrium in the environment. Formally:

Proposition 9. *For any Stackelberg-unstable game $G = (A, u_0)$, there exists an environment such that: (1) it includes only (A, u_0) and Stackelberg-stable games, and (2) it does not admit any act-rational equilibrium.*

Proof. For each Nash equilibrium payoff $v \in NE(G)$ in G, let $i(v)$ be a player and let $s_i(v)$ be a strategy such that player i can guarantee a strictly higher payoff than v_i by committing to playing mixed action $s_i(v)$ (such a strategy exists because G does not admit Stackelberg-stable equilibria). Let $\Gamma = (i(v), s_i(v))_{v \in NE(G)}$ be the set of all such pairs. Let the states in the environment be

$$\Omega = \left\{\omega_0, \left\{\omega_{(i(v),s_i(v))}\right\}_{(i(v),s_i(v))\in\Gamma}\right\}.$$

Let $u(\omega_0) = u_0$. For each state $\omega_{(i(v),s_i(v))}$, define $u\left(\omega_{(i(v),s_i(v))}\right)$ as the payoff matrix of a zero-sum game in which strategy $s_i(v)$ is the unique (possibly mixed) action that guarantees player i the value of the game $\left(A, u_{S_i(v)}\right)$. We now show that if $p(\omega_0)$ is low enough, then the environment (Ω, p, A, u) does not admit an act-rational equilibrium. Assume to the contrary that an act-rational equilibrium exists, and let v_0 be the induced Nash equilibrium in ω_0 in this act-rational equilibrium. Assume that $p(\omega_0)$ is sufficiently low. Then, it is strictly profitable for player $i(v_0)$ to deviate to partition $\left\{\left\{\omega_0, \omega_{(i(v_0),s_i(v_0))}\right\}, \{\omega\}_{\omega\in\Omega}\right\}$, and to play $s_i(v_0)$ in $\left\{\omega_0, \omega_{(i(v_0),s_i(v_0))}\right\}$ and a minimax strategy in all other states. This leads to a contradiction. □

A normal-form game is generic if all action profiles yield different payoffs. Call an environment generic if all its games are generic. The following proposition shows that any act-rational equilibrium of a generic environment is preserved when a Stackelberg-stable game is added to the environment with small enough probability. Formally:

Proposition 10. *Let $((\Omega, p), A, u)$ be a generic environment that admits an act-rational equilibrium $\tau = ((\pi_1^*, \sigma_1^*), (\pi_2^*, \sigma_2^*))$. Extend the environment by adding an additional state $\tilde{\omega}$ to Ω with probability q in which a Stackelberg-stable game $(A, u(\tilde{\omega}))$ is played, and multiply the probabilities of any state $\omega \in \Omega$ by $1-q$. Then if q is small enough, the new environment $((\Omega \cup \{\tilde{\omega}\}, \tilde{p}), A, u)$ admits an act-rational equilibrium that induces the same play in Ω.*

Proof. Consider the following meta-strategy. Each player i chooses the partition $\tilde{\pi}_i = \{\{\tilde{\omega}\}, \pi_i^*\}$. The players play a Stackelberg-stable equilibrium in $\tilde{\omega}$ with payoff v, and follow (σ_1^*, σ_2^*) in Ω both on the equilibrium path and after any player has deviated in the first stage. If player i has deviated by choosing a partition in which $\tilde{\omega}$ is part of a non-trivial partition element, then the opponent plays in $\tilde{\omega}$ a best-reply to the deviator's mixed action (which is determined by his play in the remaining states in the partition element), which yields the deviator a payoff of at most v_i (such a best reply exists due to the Stackelberg stability of $\tilde{\omega}$. The fact that $\tilde{\omega}$ τ is act-rational and the environment is generic implies that strict equilibria are played in all states. This implies that for sufficiently small q, the above strategy profile is an act-rational equilibrium. □

7 Discussion

We conclude by discussing the extension of our model to partial observability, the empirical predictions of our model, and directions for future research.

7.1 Value of Information

One can reinterpret an environment as an incomplete-information game in which each player has a minimal bound (the coarsest partition) and a maximal bound (the finest partition) on the amount of information he may acquire about the state of nature. Each player first chooses how much information to obtain within these bounds, and this choice is observable by the opponent, and then plays the incomplete information game induced by his choice. The partition profile of a rule-rational equilibrium thus describes an optimal information allocation (given his opponent's behavior): each player acquires an optimal amount of information, so that either obtaining more information or giving up information (weakly) reduces the player's payoff.

The results of Section 6 can be adapted to this interpretation as follows. Proposition 8 shows that if all games in the environment are Stackelberg-stable, then information has a positive value, and there exist an equilibrium in which both players choose to obtain the maximal amount of information. This extends related existing results for constant-sum games and games with common interests (both families of games are subsets of the set of Stackelberg-stable games). Proposition 9 shows that for any Stackelberg-unstable game, there is an environment in which it is the only game with this property, and information has a negative value: at least one of the players chooses to give up some information, and he will strictly lose by obtaining that additional information. Finally, Proposition 10 shows that if the value of information is positive at a specific equilibrium, then perturbing the environment by adding a Stackelberg-stable game with small probability maintains the positive value of information.

7.2 Partial Observability

In this section, we briefly sketch how the model and the results can be extended to a setup in which players only sometimes observe the partition of their opponent. Specifically, we change the definition of the two-stage meta-game that is induced by the environment by assuming that each player privately and independently observes his opponent's partition with probability $0 \leq q \leq 1$ (à la Dekel et al. (2007)).[12] Accordingly, we extend the definition of a meta-strategy to describe the behavior also after obtaining a non-informative signal about the opponent's partition. Finally, we define a rule-rational equilibrium as a sequential equilibrium (Kreps & Wilson, 1982) of the meta-game.

[12] The results can be extended also to a setup in which the signal about the opponent's partition is publicly observed (or has a positive probability of being observed by the opponent).

Observe that the case of $q = 0$ (i.e., non-observability of the opponent's partition) can only induce act-rational equilibria, and the case of $q = 1$ is equivalent to the basic model of Section 3. Most of the results in this paper hold also in this extended setup for all values of $0 \leq q \leq 1$: Prop. 1, 4-6, and 8-10. The remaining results (Prop. 2-3 and 9) hold for a sufficiently high $q < 1$ (the threshold may depend on the specific environment and the specific action profile).

7.3 Empirical Implications

One model implies that the predictive power of a Nash equilibrium is expected to be stronger in Stackelberg-stable games, such as zero-sum games and coordination games, relative to Stackelberg-unstable games, in which the ability to commit and to play first is advantageous. This implications fits the stylized experimental facts (see Camerer, 2003, and the references there): (1) the aggregate play in zero-sum games is close to the Nash equilibrium (see also the empirical support for this from the behavior of professional tennis players in Walker & Wooders, 2001) ; (2) players are very good at coordinating on one of the pure equilibria of a coordination game; and (3) the predictive power of Nash equilibria in most other games is more limited.

Another insight from our the model is that the way in which people play a certain game depends not only on its payoff, but also on the propensity of playing the game within the relevant environment. Moreover, our model implies that games that are played proportionally more often should yield behavior that is closer to the predictions of Nash equilibrium.

7.4 Future Research

We consider the current paper as a first step in a research agenda that studies the relations between rule rationality and environment-related commitments that operate on emotional levels (See also Winter et al., 2010). In what follows we briefly sketch two interesting directions to study in future research. First, it would be interesting to extend the model by adding small costs for using more complex playing rules (i.e., finer partitions). Second, it would be interesting to study environments that include games in extensive form (rather than normal form), and adapt the notion of rule rationality also to actions at decision nodes, and not only to strategies as a whole (e.g., a rule by which your first move is always towards the terminal node with the highest payoff).

References

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3), 488–500.

Aumann, R. J. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1), 67–96.

Aumann, R. J. (2008). Rule rationality vs. act rationality. *The Hebrew University, Center for the Study of Rationality*, DP-497.

Aumann, R. J. & Sorin, S. (1989). Cooperation and bounded recall. *Games and Economic Behavior*, 1(1), 5–39.

Azrieli, Y. (2009). Categorizing others in a large game. *Games and Economic Behavior*, 67(2), 351–362.

Bade, S., Haeringer, G., & Renou, L. (2009). Bilateral commitment. *Journal of Economic Theory*, 144(4), 1817 – 1831.

Bassan, B., Gossner, O., Scarsini, M., & Zamir, S. (2003). Positive value of information in games. *International Journal of Game Theory*, 32(1), 17–31.

Baumol, W. J. & Quandt, R. E. (1964). Rules of thumb and optimally imperfect decisions. *The American Economic Review*, 54(2), 23–46.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2), 265–272.

Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton.

Dekel, E., Ely, J. C., & Yilankaya, O. (2007). Evolution of preferences. *Review of Economic Studies*, 74(3), 685–704.

Ellison, G. & Fudenberg, D. (1993). Rules of thumb for social learning. *Journal of Political Economy*, 101, 612–643.

Fershtman, C., Judd, K. L., & Kalai, E. (1991). Observable contracts: Strategic delegation and cooperation. *International Economic Review*, 32(3), 551–559.

Frenkel, S., Heller, Y., & Teper, R. (2012). Endowment as a blessing. *Available at SSRN 2083503.*

Govindan, S. & Wilson, R. (2008). On forward induction. *Econometrica*, 77(1), 1–28.

Guth, W. & Yaari, M. (1992). Explaining reciprocal behavior in simple strategic games: An evolutionary approach. In U. Witt (Ed.), *Explaining Process and Change: Approaches to Evolutionary Economics*. University of Michigan Press, Ann Arbor.

Harsanyi, J. C. (1977). Rule utilitarianism and decision theory. *Erkenntnis*, 11(1), 25–53.

Heller, Y. (2013). Three steps ahead. *University of Oxford, mimeo.*

Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic theory*, 123(2), 81–104.

Kamien, M. I., Tauman, Y., & Zamir, S. (1990). On the value of information in a strategic conflict. *Games and Economic Behavior*, 2(2), 129–153.

Kreps, D. M. & Wilson, R. (1982). Sequential equilibria. *Econometrica*, 50(4), 863–894.

Lehrer, E., Rosenberg, D., & Shmaya, E. (2010). Signaling and mediation in games with common interests. *Games and Economic Behavior*, 68(2), 670–682.

Mailath, G. J. & Samuelson, L. (2006). Repeated games and reputations: Long-run relationships. *OUP Catalogue.*

Maynard-Smith, J. & Price, G. (1973). The logic of animal conflict. *Nature*, 246, 15.

Mengel, F. (2012). Learning across games. *Games and Economic Behavior*, 74(2), 601–619.

Mohlin, E. (2011). Optimal categorization. *Stockholm School of Economics Working Paper Series in Economics and Finance.*

Moulin, H. & Vial, J. (1978). Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3), 201–221.

Nachbar, J. H. (1990). Evolutionary selection dynamics in games: Convergence and limit properties. *International Journal of Game Theory*, 19(1), 59–89.

Osborne, M. J. & Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press, Cambrdige MA.

Renou, L. (2009). Commitment games. *Games and Economic Behavior*, 66(1), 488–505.

Samuelson, L. (2001). Analogies, adaptation, and anomalies. *Journal of Economic Theory*, 97(2), 320–366.

Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard university press, Cambridge, MA.

Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1), 25–55.

Selten, R. (1983). Evolutionary stability in extensive two-person games. *Mathematical Social Sciences*, 5(3), 269–363.

Walker, M. & Wooders, J. (2001). Minimax play at wimbledon. *The American Economic Review*, 91(5), 1521–1538.

Weibull, J. (1997). *Evolutionary Game Theory*. MIT Press, Cambridge, MA.

Winter, E., Garcia-Jurado, I., & Mendez-Naya, L. (2010). Mental equilibrium and rational emotions. *The Hebrew University, Center for the Study of Rationality, DP-521.*