RATIONALIZABILITY IN EPISTEMIC GAMES WITH ASYNCHRONOUS MESSAGES

TODD STAMBAUGH
DOCTORAL STUDENT, CUNY GRADUATE CENTER

Abstract. In 1984, Douglas Bernheim [1] and David Pearce [5] introduced the concept of rationalizable strategies in games. In the presence of common knowledge of rationality, these are the only strategies a player would ever consider. In an epistemic game, the usual definition of rationalizability provides a starting point, but it is not equipped to address the added strategic advantage a player has by virtue of having information relevant to the game, including information about other players' knowledge as noted by Rohit Parikh [2] [3]. This paper gives a definition for rationalizability in epistemic games in which knowledge is created by the sending and receiving of asynchronous messages, with varying restrictions on the types of permissible messages, using a history based approach to the logic of knowledge, along with some basic results.

1. Basic Setup

Let P be a set of propositional variables, T the set of possible truth assignments on P, $N = \{1, ..., n\}$ the set of agents, and A_i a finite set of actions for each $i \in N$. We use $A = \prod_{i=1}^{n} A_i$ to denote the set of pure strategy profiles. Utility functions $U_i : (T, A) \to \mathbb{R}$ for each $i \in N$ are required to be generic, that is each is one-to-one. We will only consider the ordering induced by the cardinal utility functions.

The goal here is to examine epistemic games in which agents initially know nothing regarding P, but learn about the truth values of propositions and other formulas through asynchronous messages. A message may be sent between players, or from nature to a player, in the sense that agents may learn about the world from external stimuli. To give an account the sending and receiving of messages, we will use a version of the history-based semantics given by Parikh and Ramanujam [4].

Let E be the set of elements of the following forms:

- (null) : nothing occurred
- $\sigma_{s,r}(\varphi)$: message φ was sent from agent s to agent r
- $\rho_{s,r}(\varphi)$: message φ was received from agent s by agent r

where $s \in N \cup \{0\}$ (0 indicating nature), $r \in N$, and φ is a formula in a language $\mathcal{L}$, which will vary depending on the types of messages admissable in each situation.

Let $\mathcal{H}_t \subset E^t$, the set of global histories of length t, be defined as follows. For any $H = (h_1 h_2 ... h_t) \in \mathcal{H}_t$, if $h_{t'} = \rho_{s,r}(\varphi)$, then it must be the case that $h_{t''} = \sigma_{s,r}(\varphi)$ for some $t'' < t'$. Clearly a message cannot be received unless it was previously sent. It need not be the case that a message that has been sent will be received and messages may be received in a different order from the order in which they were sent.

Let $w_* = (v_*, H_*)$ be the actual world and $W_t \subset T \times \mathcal{H}_t$ satisfying the following conditions:

(1) If $h_{t'} = \sigma_{0,r}(\varphi)$ for some t' then $v \models \varphi$ and $\varphi \in \mathcal{L} \cap L(P)$ and
(2) If $h_{t'} = \sigma_{i,r}(\varphi)$ for some $i \neq 0,$ then for some $k < t', 1 \leqslant j \leqslant k,$ there exist $t_j < t', s_j \neq i,$ and $\varphi_j \in \mathcal{L}$ such that $h_{t_j} = \rho_{s_j,i}(\varphi_j)$ and $\bigwedge_{m=1}^{k} \varphi_m \vdash \varphi$

Roughly, (1) indicates that nature will only send true messages about the facts of the world (not about the knowledge of agents), and (2) means that an agent will only send a true message, which must be a the logical consequences of the messages that agent has received, to ensure that the sender knows it is true.

An epistemic asynchronous message game with n players of length t is defined by $G_* = (U_1, ..., U_n, w_*)$.

We define the function $\phi_i : W_t \to E^{\leqslant t}$ for each $i \in N$ such that for $w = (v, H) \in E^{\leqslant t}$, $\phi_i(w)$ is the subsequence of events in H of the one of the forms $\sigma_{i,r}(\varphi)$ or $\rho_{s,i}(\varphi)$, where the events occur in the same order as in h. This is the local history for agent i. Define the equivalence $w \equiv_i w'$ if $\phi_i(w) = \phi_i(w')$, indicating that w and w' are indistinguishable to i. Notice that the truth assignment at w is irrelevant when it comes to the knowledge of agents, since all an agent knows about v is what is indicated by H.

2. Prerationalizability

A strategy $\beta \in A_i$ is a best response to a profile $\gamma = (\gamma_1, ..., \gamma_n) \in A$ at $(v, H) \in W_t$ if for all $\alpha \in A_i, U_i(v, (\gamma_1, ..., \gamma_{i-1}, \beta, \gamma_{i+1}, ..., \gamma_n)) \geqslant U_i(v, (\gamma_1, ..., \gamma_{i-1}, \alpha, \gamma_{i+1}, ..., \gamma_n))$. Since U_i is generic, β is uniquely defined any given γ and w.

Following the definition of rationalizability given by Pearce [5], given a game G_*, for each $i \in N$, let $A_i(0) = A_i$, and for $m = 1, 2, ...$ let $A_i(m) = \{\alpha \in A_i(m-1) \mid$ there exists $w \in W_t$ and $\gamma \in \prod_{j=1}^{n} A_j(m-1)$ such that α is a best response in $A_i(m-1)$ to γ at $w\}$.

The set of prerationalizable strategies for player i in game G_* is $Z_i(G_*) = \bigcap_{m=1}^{\infty} A_i(m)$. The set of prerationalizable strategy profiles $Z(G_*) = \prod_{i=1}^{n} Z_i(G_*)$. We use the term "prerationalizable" to acknowledge that this concept does not address the associated history, only the truth assignment. These strategies are those which are rationalizable in the case where no player knows anything about the truth assignment or the knowledge of other players. If we consider nature to be another player whose actions are defined to be the possible truth assignments on P, the prerationalizable strategies for each player are simply rationalizable in the usual sense.

3. Cases 1 & 2: Literal and Propositional Messages

We restrict our attention to the case where $\mathcal{L} = P \cup \{\neg p \mid p \in P\}$. In this situation, condition (2) on W_t can be simplified as:

$$(2)_l \text{ If } h_{t'} = \sigma_{i,r}(\varphi) \text{ for } i \neq 0, \text{ then for some } t'' < t' \text{ and } s \neq i, h_{t''} = \rho_{s,i}(\varphi)$$

Under this framework, each agent can only have knowledge of literals and the knowledge of literals of agents from whom she received a message. That is, if r receives message φ from s, then r will know that φ is true, but also that s knows that φ is true, since the message could not have been sent otherwise. In later cases, we will update condition (2) to reflect this reasoning, but for now it is not necessary.

Define $Z_{j,i}(G_*) = \{\alpha \in Z_j(G_*) \mid$ There exist $w, w_1 \in W_t, \gamma \in Z(G_*)$ such that α is a best response for j to γ at w and $w \equiv_j w_1 \equiv_i w_*\}$. This is the set of strategies which are best responses for j as far as i can tell. The actual set of strategies j will consider may be smaller, since j will only consider the possibility that $\phi_j(w) = \phi_j(w_*)$, a more restrictive condition. As such, i will respond to j-strategies which may not be rationalizable, since j may have more information than i can tell.

It is worth noting that while j may have learned information from messages from other agents, and indeed this will be the case for some choices of w, i will always need to consider possible options for w in which j only has knowledge obtained directly from nature, and thus knows nothing about other agents' knowledge. This is the primary reason that i has to consider j-strategies that respond to all possible best responses on the part of other players.

The set of rationalizable strategies for i is defined as follows:

$R_i(G_*) = \{\alpha \in Z_i \mid$ there exist $w \in W_t$ and $\gamma \in \prod_{j=1}^{n} Z_{j,i}$ such that $w =_i w_*$ and α is a best response at w to $\gamma\}$

When the game is clear from context we will omit the reference to G_*, e.g. $R_i(G_*) = R_i$.

If we extend $\mathcal{L}$ to $L(P)$, little changes. In fact, the set of rationalizable strategies can be defined in exactly the same way. The only difference of note is that we now must use the original condition (2).

4. Epistemic Messages

So far we have considered only histories involving messages in the language $L(P)$, so before we can consider the implications of epistemic messages on rationalizability, we must expand the framework to include such messages, with the appropriate restrictions. We define the language of epistemic games with n players on P, denoted $EGL_n(P)$, as follows.

(i) If $\varphi \in L(P)$, then $\varphi \in EGL_n(P)$.
(ii) If $\varphi, \psi \in EGL_n(P)$, then so are $\neg\varphi$ and $\varphi \wedge \psi$.

(iii) If $\varphi \in EGL_n(P)$, then so is $K_i\varphi$ for each $i = 1, ..., n$.

As in [4], the semantics for $EGL_n(P))$ are given as follows.

(i) If $\varphi \in L(P)$, then $(v, H) \models \varphi$ iff $v \models \varphi$.
(ii) If $(v, H) \models \neg\varphi$ and $(v, H) \models \varphi \wedge \psi$ are defined in the usual way.
(iii) $w = (v, H) \models K_i\varphi$ iff for all $w' = (v', H')$ with $w \equiv_i w'$, $w' \models \varphi$.

In addition to the new language, we must expand our notion of what messages are permissible to send. As before, we require that agents send only messages which they know are true, but the new expressive power provided by $EGL_n(P)$ adds an additional wrinkle. Since the structure of the game is common knowledge, including this requirement, an agent receving message φ from some other agent j, learns not only φ, but $K_j\varphi$. This was true in previous cases, but since messages were only contained in $L(P)$, the receiver could act on $K_j\varphi$ only in terms of strategy. To reflect this shift, we must update condition (2) in the following way.

$(2)_K$ If $h_{t'} = \sigma_{i,r}(\varphi)$ for some $i \neq 0,$ then for some $k < t', 1 \leqslant j \leqslant k,$ there exist $t_j < t', s_j \neq i,$ and $\varphi_j \in L_K(P)$ such that $h_{t_j} = \rho_{s_j,i}(\varphi_j)$ and $\bigwedge_{m=1}^{k} K_{s_m}\varphi_m \vdash \varphi$.

An agent can now send messages following not only from received messages, but from the fact that the senders of those messages knew them to be true.

5. Case 3: Messages Involving One Epistemic Operator

Before tackling all of $EGL_n(P)$, we will first consider the simpler case where $\mathcal{L} = L(P) \cup \{K_j\varphi \mid \varphi \in L(P) \text{ and } j \in N\}$. This restriction is purely artificial, but will inform the process moving forward. In the previous cases, agents were able to reason about the strategies of other players based on what each of those players knew about the world, but this reasoning only went one level up. As far as i knew, j had to respond to all of the prerationalizable strategies of all of the other players.

Upon receiving a message $K_{j_1}\varphi$ from j_2, agent i will know the formula $K_{j_2}K_{j_1}\varphi$. As such, i will be able to strategize based on the same thinking described in the previous section for both j_1 and j_2, but additionally, i will have additional information about which strategies of j_1 will be considered by j_2 by attributing to j_2 the same ability to reason about $K_{j_1}\varphi$ as i would employ. To accommodate this type of reasoning, we will buld upon the previous definition in a similar fashion.

Define $Z_{j_1,j_2,i} = \{\alpha \in Z_{j_1} \mid$ There exist $w, w_1, w_2 \in W_t, \gamma \in Z$ such that α is a best response for j_1 to γ at w and $w \equiv_{j_1} w_1 \equiv_{j_2} w_2 \equiv_i w_*\}$. As before, this is the set of best responses for j_1 considered possible to j_2, as far as i can tell. Because of the limitations still in place on messages, this is as far out as i will need to consider, just one epistemic level beyond what was in play previously.

Since some of the actions in Z_{j_1} may be eliminated at this level, i will need to assess strategies for j_2 with this in mind. As such, it is necessary to redefine $Z_{j,i}$ as:
$Z_{j,i} = \{\alpha \in Z_j \mid$ There exist $w, w_2 \in W_t, \gamma \in \prod_{k=1}^{n} Z_{k,j,i}$ such that α is a best response

for j to γ at w and $w \equiv_j w_2 \equiv_i w_*\}$. Note that the only change from before is that now γ will be taken from the set of strategy profiles that i deems possbile for j to consider. If no messages including epistemic operators have been sent from j to i, then $\prod_{k=1}^{n} Z_{k,j,i} = Z$, so this new definition will agree with the previous one under the condition that messages contain no epistemic operators.

All of the added complexity of this case arises at the levels of reasoning described already, so the set of rationalizable strategies is defined exactly as before, that is $R_i = \{\alpha \in Z_i \mid$ there exist $w \in W_t$ and $\gamma \in \prod_{j=1}^{n} Z_{j,i}$ such that $w =_i w_*$ and α is a best response at w to $\gamma\}$.

6. The Role of Negations of Epistemic Operators

One important step in defining rationalizability$EGL_n(P)$ is to address the interaction between connectives and epistemic operators. Of fundamental importance is the fact that with communication, negations of knowledge formulas can change truth values over time. This phenomenon exists in any epistimic situation where communication is done over time, and even under certain conditions where communication occurs at a single time.

We have used the fact that agents may only send messages they know to be true, and this will remain true when we expand the language to include negations of epistemic formulas. Regardless of the language, conditions (1) and $(2)_K$ naturally form an induction argument that only true messages are sent. Clearly, by (1), nature will only send true messages. In a history H, if we assume that all messages sent are true prior to time t', and $h_{t'+1} = \sigma_{s,r}(\varphi)$, then by $(2)_K$, φ is a consequence of the knowledge of the senders of previous messages, and since each of those messages must be true and the senders must have known that the messages were true, then φ itself must be true.

If it were the case that φ was such a formula that was true at the time the message was sent, but later became false, in general the sender of that message would not be able to tell at what time the message became false, since the messages are asynchronous, and the sender is only aware of the messages she sends and receives. This means that in this situation, φ may have become false between her reception of the messages that she used to deduce φ and her sending the message as far as she can tell. Since φ must be true at the time of her sending the message, this turn of events is impossible, so φ can never become false.

While there are formulas that may change truth values in an epistemic game, the structure here simply precludes such messages from being sent. We will leave questions regarding this phenomenon for future work.

7. Case 4: Messages in $EGL_n(P)$

Finally, we follow same method for generalizing the language as the previous case to address $\mathcal{L} = EGL_n(P)$. As discussed above, negations present no new complications, and unsurprisingly neither do conjunctions, so the final issue to address is that of nested epistemic operators. Now we have no restrictions on the length and complexity of messages, so we cannot rely on the language to dictate how many levels out the players will need to reason.

Since the histories we consider have length t, we do have a natural limit to the epistemic complexity that messages can take, i.e. x has length at most $(t-1)$ (since the first message cannot contain any epistemic operators). Obviously, with the restrictions in place already, we can lower this bound by at least half, but for the purposes here, any bound will do. We could even loosen the condition on common knowledge of the structure of the game to include the possibility that there is common knowledge that each player has an upper bound for t, which may be completely private. In that situation, each player could use that upper bound as a starting point. For simplicity we will maintain common knowledge of t. This gives the starting point for the agents' reasoning, defined as follows for each $j_1, ..., j_t, i \in N$.

$Z_{j_1,...,j_t,i} = \{\alpha \in Z_{j_1} \mid$ There exist $w, w_1, ..., w_t \in W_t, \gamma \in Z$ such that α is a best response for j_1 to γ at w and $w \equiv_{j_1} w_1 \equiv_{j_2} ... \equiv_{j_t} w_t \equiv_i w_*\}$.

From there the agents will reasoning will proceed in the same way as before, so for each $t' < t$ and $j_1, ..., j_{t'}, i \in N$:

$Z_{j_1,...,j_{t'},i} = \{\alpha \in Z_{j_1} \mid$ There exist $w, w_1, ..., w_{t'} \in W_t, \gamma \in \prod_{k=1}^n Z_{k,j_1,...,j_{t'},i}$ such that α is a best response for j_1 to γ at w and $w \equiv_{j_1} w_1 \equiv_{j_2} ... \equiv_{j_{t'}} w_{t'} \equiv_i w_*\}$.

and

$R_i = \{\alpha \in Z_i \mid$ there exist $w \in W_t$ and $\gamma \in \prod_{j=1}^n Z_{j,i}$ such that $w =_i w_*$ and α is a best response at w to $\gamma\}$.

8. Basic Results

For a game of length $(t-1)$ given by $(U_1, ..., U_n, (v, H'))$, it is natural to generate a game of length t given by $(U_1, ..., U_n, (v, H))$ simply by amending H' with some event h_t, according to the restrictions outlined in the first section.

Lemma (1)**.** *If $w = (v, H'h_t)$ for some $h_t \in E$ and $w \equiv_i w_1$, then there exist some H_1', H_1'' such that $w_1 \equiv_i (v_1, H_1'h_tH_1'')$ and $w' \equiv_i w_1'$ where $w_1' = (v_1, H_1'H_1'')$.*

Proof. If $h_t = \sigma_{i,r}(\varphi)$ or $h_t = \rho_{s,i}(\varphi)$, then since $w \equiv_i w_1$, h_t must appear somewhere in H_1, so we simply choose H_1' and H_1'' such that $w_1 = (v_1, H_1'h_tH_1'')$.

If $h_t \neq \sigma_{i,r}(\varphi)$ and $h_t \neq \rho_{s,i}(\varphi)$, then $\phi_i(w') = \phi_i(w) = \phi_i(w_1)$, trivially, since h_t will not be visible under ϕ_i. Let H_1' be an initial substring of H_1 which is followed by some e such that $e \neq \sigma_{i,r}(\varphi)$ and $h_t \neq \rho_{s,i}(\varphi)$ which is then followed by H_1''. Such a choice must exist since the length of $\phi_i(w_1)$ is less than the length of H_1. Since H_1' and H_1'' are chosen by deleting an event from H_1 not visible to ϕ_i, clearly $\phi_i(w_1) = \phi_i(w_1')$. □

Lemma (2). *If* $w \equiv_i (v_1, H_1' h_t H_1'')$*, then* $(v_1, H_1' h_t H_1'') \equiv_i (v_1, H_1' H_1'' h_t)$*.*

Proof. If $h_t \neq \sigma_{i,r}(\varphi)$ and $h_t \neq \rho_{s,i}(\varphi)$, h_t is not visible to i, so its placement in the global history has no effect on ϕ_i. The conclusion follows trivially.

If $h_t = \sigma_{i,r}(\varphi)$ or $h_t = \rho_{s,i}(\varphi)$, then H_1'' cannot contain any events of either of those forms, otherwise the hypothesis would fail. For the same reason as the previous case, this means movement of the events in H_1'' within the history does not change the evaluation under ϕ_i, so again the conclusion is easily shown to be true. □

One simple intuition about rationalizability in epistemic games is that the set of a player's rationalizable strategies should shrink, or at least not grow upon learning new information. The following theorem shows not only this, but that the set of rationalizable strategies for every player will not grow upon additional communication.

Theorem. *If* $H_* = H_*' h_t$ *for some* $h_t \in E$ *such that* $G' = (U_1, ..., U_n, (v_*, H_*'))$ *and* $G = (U_1, ..., U_n, (v_*, H_*))$ *are epistemic asynchronous message games of length* $(t-1)$ *and* t *respectively, then* $R_i(G) \subseteq R_i(G')$*.*

Proof. We will show that at every stage of reasoning, the set of strategies under consideration in G is a subset of those in G'. Since the prerationalizable strategies do not take histories into consideration, clearly $Z_i(G') = Z_i(G)$ for each i.

As discussed previously, if t exceeds the possible epistemic complexity of messages in the game, then $Z_{j_1,...,j_t,i} = Z_{j_1}$. Since G' has length $(t-1)$, it must be that $Z_{j_1,...,j_t,i}(G') = Z_{j_1}(G')$. By definition, $Z_{j_1,...,j_t,i}(G) \subseteq Z_{j_1}(G)$. This means that $Z_{j_1,...,j_t,i}(G) \subseteq Z_{j_1,...,j_t,i}(G')$ for any $j_1, ..., j_t, i$.

Assume $Z_{j_1,...,j_{t'+1},i}(G) \subseteq Z_{j_1,...,j_{t'+1},i}(G')$ for some t' and any $j_1, ..., j_{t'+1}, i$ and suppose $\alpha \in Z_{j_1,...,j_{t'},i}(G)$. This means there must be some $w, w_1, ..., w_{t'} \in W_t$ and $\gamma \in \prod_{k=1}^n Z_{k,j_1,...,j_{t'},i}(G)$ such that α is a best response for j_1 at w to γ and $w \equiv_{j_1} w_1 \equiv_{j_2} ... \equiv_{j_{t'}} w_{t'} \equiv_i w_*$.

By assumption we have that $\prod_{k=1}^n Z_{k,j_1,...,j_{t'},i}(G) \subset \prod_{k=1}^n Z_{k,j_1,...,j_{t'},i}(G')$, so the same choice of γ may be made in G'.

Using the lemmas above, and the choice given for $w, w_1, ...w_{t'}$, it must be the case that $w_* \equiv_i w_{t'} \equiv_i (v_{t'}, H_{t'}' h_t H_{t'}'') \equiv_i (v_{t'}, H_{t'}' H_{t'}'' h_t)$ for some choice of $H_{t'}'$, $H_{t'}''$ such that $w_*' \equiv_i w_{t'}'$. We repeat this process using the fact that $w_{t'} \equiv_{j_{t'}} w_{t'-1}$ and $w_{t'}$ is generated from $w_{t'}'$ in the same way as w_* from w_*'. This gives a chain of worlds in $W_{(t-1)}$, $w' \equiv_{j_1} w_1' \equiv_{j_2} ... \equiv_{j_{t'}} w_{t'}' \equiv_i w_*'$. Since the truth assignment in w' is the same as that of w', utilities will match, so α is a best response for j_1 to γ at w' as well.

This means that $\alpha \in Z_{j_1,...,j_{t'},i}(G')$, and therefore $Z_{j_1,...,j_{t'},i}(G) \subseteq Z_{j_1,...,j_{t'},i}(G')$.

Since R_i is constructed by the same process as each of the $Z_{j_1,...,j_{t'},i}$, by applying the previous argument inductively, $R_i(G) \subseteq R_i(G')$.

□

Since no strategy that was not rationalizable will become so upon further communication, we can see that the rationalizable strategies will stabilize not only for finite games, but for games with infinite histories. Define $\mathcal{H}_\infty$ to be the set of global histories $H = (h_1 h_2, ...)$ with the same constraints as those for the finite case. We can similarly define $w_* \in T \times \mathcal{H}_\infty$ and an infinite epistemic game $G_* = (U_1, ..., U_n, w_*)$. We will now let $R_i(G_*, t)$ be the rationalizable strategies for the game with history of length t given by the initial subsequence of H_* with length t and the following corollary defines $R_i(G_*)$ for the infinite case.

Corollary. *For the infinite epistemic game G_*, there exists some t such that for all $t' > t$, $R_i(G_*, t) = R_i(G_*, t')$.*

Proof. The set of strategies A_i is finite, so $R_i(G_*, 0)$ must also be finite. By the preceding theorem, $R_i(G_*, t) \subseteq R_i(G_*, t-1)$. This forms a chain of subsets beginning with one of finite cardinality, so the size of subsequent sets may only go down as many as $\mid R_i(G_*, 0) \mid$ times, and one of those times must be the last.

□

9. Conclusion

A history based approach to the study of epsitemic games provides a strong framework for examining reasoning at every individual level of knowledge, as opposed to all at once. The asynchronous message games discussed here are one of many kinds of games that this approach may illuminate. Future work may take a similar line to discuss games with different kinds of messages, as well as different mechanisms by which those messages are sent and received. Further, it would be worth examining other solution concepts beyond rationalizability using a similar approach. Questions of complexity of computation and computability in the case of infinite games remain to be answered, in order to examine not only the robustness of the definition given, but also its practicality for use in modeling various strategic epistemic scenarios. Whatever path one chooses from here, the topic of epistemic games remains rich territory for future work.

References

[1] Bernheim, B. Douglas, "Rationalizable Strategic Behavior," *Econometrica*, Vol. 52, No. 4 (July, 1984), 1007-1028

[2] Parikh, Rohit "An Epistemic Generalization of Rationalizability," Submitted (2016)

[3] Parikh, Rohit "Knowledge and Action in Groups," *Studies in Logic*, Vol. 8, No. 4 (December, 2015), 108-123

[4] Parikh, Rohit; Ramanujam, Ramaswamy, "A Knowledge Based Semantics of Messages," *Journal of Logic, Language and Information*, 12.4 (2003): 453-467

[5] Pearce, David G., "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, Vol. 52 No. 4 (July, 1984), 1029-1050