Awareness Equilibrium

Jernej Čopič and Andrea Galeotti *

This Draft: November 14, 2005

Abstract

We define awareness of rational agents by their perceptions. Awareness Equilibrium is then given by agents' perceptions of the world, their conjectures about others' perceptions, actions that are taken, and outcomes consistent with all of these. We first consider economic problems where agents may lack awareness of other agents. These setups have a natural interpretation in terms of social networks. We propose a model of technological choice with no payoff externalities. An example akin to this model is viral marketing. Second, we apply Awareness Equilibrium to strategic situations (normal-form games) where players may be aware only of subsets of actions. Generically the set of Action-awareness Equilibrium outcomes is much larger than the set of Nash equilibria. This provides a critique of the approach pursued by the epistemic literature on *unawareness*. We outline an evolutionary approach for the study of awareness.

1 Introduction

Awareness defines the way that individuals perceive the world. Lack of awareness is a lack of cognition which may stem from the lack of attention, consciousness, or bounded intellectual resources. All of these are inherently unobservable parts of an individual, and from an economic perspective, what ultimately matters is behavior. Our approach is to define awareness as an equilibrium notion where conjectures not only have to be internally consistent and consistent with agents' rationality, but also have to be consistent with the outcome that actually obtains. *Awareness Equilibrium* is thus given by agents' perceptions of the world, their conjectures about others'

*jernej@hss.caltech.edu, agaleotti@ist.caltech.edu. A conversation with Leeat Yariv helped us to get started on this project. We have also benefitted from discussions with Ken Binmore, Colin Camerer, Matt Jackson, Andrea Mattozzi, Debrah Meloso, Bill Zame, and the participants at the BWED conference 2005, especially Beth Allen, Andres Hervas and Tural Huseynov, and the Caltech SISL 2005 workshop.

perceptions, actions that are taken, and outcomes consistent with all of these. On the one hand, the outcome that obtains is revealing of the agents' awareness, and on the other hand there may be many awareness structures that are consistent with a particular outcome.

Consider a simple economic example where awareness matters. To illustrate how Awareness Equilibrium works we first give the outcome and the equilibrium as perceived by the agents, and then we give the whole situation as it would have been perceived by an omniscient outside observer. Two Slovenian tourists coincide for two nights in a hotel in New York City. When they meet the first evening they both decide to buy a digital camera the next day. They discuss the quality of various brands, and each recalls being recommended the specific Brand, by some perfectly random person he met that day. Each forms a conjecture that Brand is good, and the next day independently makes the purchase. When they meet again the second evening, they observe that the other has purchased the Brand as well. This outcome confirms their conjectures on the quality of Brand.

What actually happened is as follows. A random Italian tourist was standing at a traffic light. He asked a Slovenian to take a picture of him with Italian's own brand new Brand camera, after which the Italian excitedly showed and explained all the fantastic Brand's features. This being a perfectly usual random incident, none of the two Slovenians payed much attention to the specifics. They were not *aware* that the Italian tourist had been payed by Brand to stand on that corner and advertise. Everything that Slovenians observed and *reasoned* is *consistent*, given their lack of awareness. This marketing strategy is an example of *viral marketing*, a persuasive form of advertising that is used extensively.

In this example, even though Slovenians are perfectly rational and consistent thinkers, none of the two can within the model figure out the mistake that the information of the Italian is a random unbiased observation. The observed outcome along with such awareness and conjectures is an equilibrium, which we call Awareness Equilibrium.

In Section 2 we define awareness as an equilibrium concept as follows. Agents are rational but may be aware of only a part of their environment, and they make conjectures about awareness of others. Agents' conjectures are consistent so that if an agent is not aware of something he may not reason about it. Agents' pick their actions optimally, given their awareness and conjectures. Finally, when the outcome is realized each agent must see that outcome being consistent with his awareness, conjectures and rationality. Thus in Awareness Equilibrium, a person cannot become aware of new things. That is what an "equilibrium" entails; for if things must change, then it is not really an equilibrium. Importantly, Awareness Equilibrium ties unobservable perceptions to the observable outcome.

Such revealed awareness is conceptually the core of our approach and is consistent with other notions of equilibrium. Awareness Equilibrium is close to the equilibrium models where off the equilibrium deviations are only conjectures but never actually observed (see Rubinstein and Wolinsky [1991], Fudenberg and Levine [1993], and Battigalli and Guaitolli [1998]). That is a very natural weakening of standard notions of equilibrium, and our Awareness Equilibrium weakens it further by not requiring that the model itself be common knowledge. The present notion entails complete rationality and bounded cognition, but it can easily be extended to account for bounded rationality as well.

In Section 3 we study economic problems where agents have limited awareness of others. We provide a formal mapping from awareness of agents, and conjectures thereof, to a collection of networks. We call this construction the *social network*. For example, the complete social network characterizes common knowledge of agents. Congruently with our approach we ask whether an observed outcome reveals the underlying social network. As an application, we provide a new model of choice of technology with no payoff externalities. There are two technologies, one probabilistic and the other deterministic. Agents sample the probabilistic technology. Each agent infers the quality of the probabilistic technology using his sample and the conjectured samples of others. Conjectures are about the awareness of others and about consistent mappings of their awareness and samples into their choices. Based on their inferences, agents choose technologies according to a rationalizability criterion. Agents observe choices of others and in equilibrium these choices have to be self-confirming with respect to agents' conjectures.

In the framework of Awareness Equilibrium the social network has an impact on technological choice. We show that there is a unique equilibrium where agents choose the deterministic technology regardless of their sample, even when that is inefficient. In that *misperception* equilibrium each agent incorrectly, albeit consistently, conjectures that the other agents are not aware of each other, but they are only aware of him. This causes the agents to effectively misperceive the correlation between the actions of the other two, and consequently over-rate the informational content of their actions, relative to own sample. Observing such outcome reveals the social network of misperception.

In Section 4 we apply the Awareness Equilibrium to normal-form games. There the set of players is common knowledge but they may be aware only of a part of the action set. We interpret the level of awareness of agents, i.e. how many actions they are aware of, as the agents' cognitive bound. The study of equilibrium outcomes for different levels of cognitive bounds provides a specific econometric method to measure awareness from observed outcomes. We demonstrate that generically the set of Action-awareness Equilibrium outcomes is a much larger superset of the set of Nash Equilibria. As the agents' cognitive bound increases, the former set shrinks towards the latter. Nevertheless, there do not exist restrictions on games which would make these two sets coincide for a given cognitive bound while allowing for the games to have an arbitrary size of action sets.[1]

These results provide a critique of the focus of the existing literature on epistemic foundations of *unawareness*. That literature provides axiomatic constructions of awareness and conjectures, e.g. Feinberg [2005], Heifetz et al [2004], Li [2004], Modica and Rustichini [1999], Dekkel et al [1998]. In applications, the literature focuses entirely on games and awareness of actions and it is always the case that awareness constructions are exogenously imposed on agents. In contrast with our equilibrium approach the literature does not require that observed outcomes must be consistent with awareness. Thus, Action-awareness Equilibrium outcomes always form a strict subset of outcomes under notions of Nash Equilibrium adapted to epistemic awareness structures proposed by the literature, *when all admissible awareness constructions are allowed*.[2] The critique is therefore two-fold. On one hand, the approach of the epistemic literature on unawareness is ad-hoc in the sense that awareness is imposed on agents exogenously. On the other hand, if awareness were not imposed exogenously, those notions of unawareness immediately lead to "everything goes" results, since consistency with observables is not required.

Finally, in Subsection 4.3 we outline a new research agenda for the study of awareness in an evolutionary framework. That is an evolutionary model of species with limited resources which Nature allocates among various facets of intelligence. One of these facets is the quantifiable cognitive bound of the species. The bounded cognition model of the Action-awareness Equilibrium is a pillar of that approach. In Section 5 we conclude.

2 Awareness Equilibrium

A problem is described by some universal set $\mathcal{U} = \mathcal{U}_{ob} \cup \mathcal{U}_{con}$. In $\mathcal{U}_{ob}$ there are all the objective facts, for example the agents, $\mathbf{N} \subset \mathcal{U}_{ob}$, their action sets $\mathbf{A}_n$, $n \in \mathbf{N}$, exogenous facts, $\mathbf{F}$, etc. In $\mathcal{U}_{con}$ there are conjectures that agents can make about the world around them, more precisely about subsets of $\mathcal{U}$. We define *awareness* as equilibrium notion. See Appendix A for a formal definition.

Definition **1.** Awareness Equilibrium (AE) is defined as follows.

AE1 Each agent $n \in \mathbf{N}$ is *aware* only of a subset of the universal set, $\emptyset \neq \mathcal{U}_n \subset \mathcal{U}$. He makes conjectures about awareness of every agent he is aware of, and in

[1] Agents' cognitive bounds describe the limitation to their attention imposed by Nature, and should thus be fairly constant across different games. In particular, it doesn't make much sense to *a priori* assume that in a larger game, agents should have a much larger cognitive bound.

[2] Admissible awareness constructions are those that satisfy a given set of axioms. We discuss this in more depth in Subsection 4.2.

general about an element of $\mathcal{U}$ if and only if he is aware of it. *Each agent understands that this is how things are done.*[3]

AE2 Each agent's conjectures are consistent: for instance, if agent n conjectures that agent k is aware of $\alpha \in \mathcal{U}$ then n has to be aware of α and k (same at higher orders); if n is aware of himself then the conjecture about own awareness is correct (same at higher orders). *Each agent understands that this is how things are done.*

AE3 Given $\mathcal{U}_n$, each agent n picks from his optimal actions in $\mathcal{U}_n$. *Each agent understands that this is how things are done.*

AE4 Let $o \in \mathcal{O}$ be the outcome corresponding to agents' choices of actions and let o_n be the restriction of o, observable to n given $\mathcal{U}_n$. Then o_n is consistent with n's conjectures and *understanding.*

The first three points describe how agents perceive and approach the world around them. More precisely, requirement AE1 says that agents do not need to be aware of the whole universe, and it describes what the agents are allowed to reason about. Specifically, agents are prohibited from reasoning about facts of which they are not aware, and that never reasoning about a fact is equivalent to not being aware of it. Also, an agent has to be aware of something, since otherwise he would in equilibrium have no impact on the economy so he would be irrelevant. Requirement AE2 says that agents have to make their conjectures consistently with the facts they are aware of, and if an agent is aware of himself then he is aware of his own conjectures. This applies to all orders of conjectures, so that for example, if an agent conjectures a specific conjecture of another agent, he must conjecture that the other agent is aware of his own conjecture. Requirement AE3 says that agents are rational in their choice of actions but limited by their awareness.

Requirement AE4 ties awareness to the outcome in a way consistent with the agents' reasoning. That is, if an outcome occurred, which was perceived to be inconsistent with agents awareness, conjectures, or rationality, then rational agents would figure out that something was wrong with their perception of the world. The mapping from actions into outcomes is most often the standard one. For instance, in a "game" an outcome is simply a vector of payoffs, arising from actions that are chosen. AE4 requires that only the part of the outcome observable to n under his awareness should be consistent with his awareness and conjectures. Thus, if in a "game" we assume that agent n only observes payoffs to agents he is aware of, then

[3]When an agent conjectures reasoning of another agent that he is aware of, he correctly conjectures that the reasoning of that other agent behaves according to the given axiom. Thus, this phrase is a common-knowledge-like restriction on higher order conjectures. We use the phrase because saying "This is common knowledge" could cause confusion - for example, it is not common knowledge to what set of agents the axiom applies.

from the point of view of n only the payoffs to those agents have to be in equilibrium consistent with n's conjectures and rationality.

Observe that we allow in principle (in terms of conjectures) for an agent to be aware only of elements of $\mathcal{U}_{con}$. However, in an AE, every agent has to be aware of some objective fact, as long as the outcome function is non-trivial. The following example illustrates this, and more generally, it illustrates how AE describes awareness.

Example **1.** This example addresses the following question: "Why is there something and not nothing?" The following is a simple setup addressing this question in our framework. There is one agent, $\mathbf{N} = \{1\}$ and $\mathcal{U}_{ob} = \mathbf{N}$. The outcome function is defined as the mapping $O : \{\mathcal{U}\} \to \{0, 1\}$, where $O(\mathcal{U}) = 1 \iff \mathcal{U} \neq \emptyset$. In most situations it makes sense to assume that $O \in \mathcal{U}_{ob}$, but for the purpose of this example we assume that $O \in \mathcal{U}_{con} \cap \mathcal{U}_1$. Agent 1 can make conjectures. In particular, 1 can make the following three conjectures:

$$C1 \equiv (\mathcal{U} \neq \emptyset), C2 \equiv (\mathcal{U}_{ob} = \emptyset), C3 \equiv (\mathcal{U} = \emptyset)$$

Clearly, $C3$ cannot be a part of an AE, since it violates the consistency in AE2: $C3 \in \mathcal{U}$, hence $C3$ is not correct. Second, $C2$ can also not be a part of an AE. That is, $C2$ does not violate AE1-AE3, but from the 1's perspective $C2$ is inconsistent with the outcome (which is 1) since if there are conjectures *somebody* must be making them, but $C2$ asserts that there is no one. Conjecture $C1$ may form a part of an AE. Specifically, $C1a \equiv [(\mathcal{U}_{ob} \neq \emptyset) and (\mathcal{U}_{con} \neq \emptyset)]$ can be a part of an AE. Our answer to the question is therefore that there is something because someone is asking the question, so that at the least there is someone and there is the question.

Some aspects of the problem, a subset of $\mathcal{U}_{ob}$, may be common knowledge. An easy way to incorporate this into our definition is as specific restrictions on conjectures. For example, in a two-agent strategic interaction it usually makes sense that the two players are aware of each other and that is common knowledge. We assume that utility functions over all outcomes are given, i.e. $u_n : \mathcal{O} \to I\!R$. Consistently with the equilibrium notion, agent i is only *aware* of the restriction of $u_n : \mathcal{O}_n \to I\!R$, where $\mathcal{O}_n = \mathcal{U}_n \cap \mathcal{O}$ is the subset of outcomes that n is aware of. Similarly for n's conjectures about other agents' utility functions. We assume in all our examples that the mapping from actions and facts into outcomes is common knowledge (again, modulo awareness, as with utility functions). In any case, which aspects are common knowledge, and which are not, is a matter of the specific context.

The notion of awareness we propose simply *is* the equilibrium concept where awareness is invariably linked to the outcome. This is necessary if awareness is to be a useful concept, since awareness occurs in agents' minds and is therefore not something measurable *per se*. On one hand by observing an outcome, one can establish the awareness that supports it as an AE. On the other hand, by asserting that agents must be aware of subsets of $\mathcal{U}_{ob}$ with, for example, some specified cardinality, some outcomes can be ruled out. In the rest of this paper, we provide a few specific frameworks in which we apply the definition of Awareness Equilibrium.

3 Social-awareness Equilibrium

In this Section we consider situations where agents may not be fully aware of other agents. Throughout this Section we assume for convenience that each agent is aware of himself.

Congruent with our approach is the idea that by observing an outcome in a society we can infer something about the social awareness of the agents, which have a very natural interpretation in terms of network-like objects. In this spirit, we define the *Social-awareness Equilibrium* SAE. First, in Subsection 3.1., we interpret awareness and conjectures as a collection of networks, i.e. we provide a formal mapping from awareness to networks. In Subsection 3.2 we provide a specific definition of SAE, adapted to our example of R&D. There the agents' social awareness affects the choice of technology and may lead to very inefficient outcomes. We call this an example of "misperception".

3.1 Mapping Awareness to Networks

Each agent is aware of a subset of agents, and makes conjectures about awareness of others. We denote by $\mathbf{N}^{(n)}$ the set of agents that n is aware of, and by $\mathbf{N}^{(n,m)}$ the conjecture of n about m's awareness and so on. Let

$$CN_n = (\mathbf{N}^{(n)}, \mathbf{N}^{(n,m)}, ...).$$

Now we interpret CN_n as a collection of networks. Let

$$g^{(0)} : \times_{n\in\mathbf{N}}\mathbf{N}^{(n)} \rightarrow \{0,1\}^{|\mathbf{N}|^2},$$

such that $g^{(0)}_{n,m} = 1$ if and only if $m \in \mathbf{N}^{(m)}$. The *directed links* in $g^{(0)}$ describe awareness. Let

$$g^{(n)} : \times_{m\in\mathbf{N}^{(n)}}\mathbf{N}^{(n,m)} \rightarrow \{0,1\}^{|\mathbf{N}|^2},$$

such that $g^{(n)}_{m,l} = 1$ if and only if $l \in \mathbf{N}^{(n,m)}$, $\forall n \in \mathbf{N}$. The directed links in $g^{(n)}$ describe first order conjecture of awareness of agent n. In analogous fashion, we define the whole collection of networks, $G = (g^{(0)}; g^{(n)}; g^{(n,m)}; g^{(n,m,l)}....), n, m, l, ... \in \mathbf{N}$, which is a one-to-one representation of agents' awareness and awareness conjectures. That is, each such collection of networks represents a unique $CN = (CN_n), n \in \mathbf{N}$ and conversely each CN has unique such representation G. We can call G the *social network*.

When $g^{(n,m,l,...)}$ is a complete network (i.e. $g^{(n,m,l,...)}$ maps onto $\{1\}$) for all finite sequences $(n, m, l, ...)$, $n, m, l, ... \in \mathbf{N}$, we call G the *complete social network*. The following remark is immediate, and it illustrates how the awareness of agents relaxes common knowledge.

REMARK **1.** G is a complete social network if and only if the set of agents $\mathbf{N}$ is common knowledge.

There are *a-priori* no restrictions on G. However, AE imposes consistency restrictions on G, e.g. in every AE, $n \in \mathbf{N}^{(n)}$ and $\mathbf{N}^{(n,n)} = \mathbf{N}^{(n)}$. Moreover, in an AE different consistent G's will sustain different equilibrium outcomes. Thus, we can ask the question - Does an observed outcome *reveal* the underlying social network? With this, we turn to the example of technological choice.

3.2 Misperception and Technology Choice

In this subsection we study a model of choice of technology. In this model there are no payoff externalities, but the awareness and conjectures have an impact on AE outcomes. We define the SAE in terms of CN, which by above is equivalent to defining it in terms of G. We then interpret our results in terms of G.

There are 3 agents, $\mathbf{N} = \{1, 2, 3\}$, choosing between two different technologies. The payoff from one technology is deterministic, α, while the payoff to the other technology is probabilistic, $\beta \in \{\beta_L, \beta_H\}$, where $\beta_L < \alpha < \beta_H$. Let ν be the true probability of β_H, $\nu \in (0, 1]$ but the agents have no knowledge or prior over ν. There are no payoff externalities, and the agents are risk neutral. We will later further specify the relationship between α, β_L, β_H.

Each agent samples one draw from β. We stress that our results does not depend on the fact that agents observe one sample of β; what is important is that they observe a finite sample. We consider 1-sampling for analytical simplicity. It is also clear that our method can be applied to an arbitrary finite number of players, and that the results are then qualitatively the same.

The action set of each agent is $A_n = A = \{\alpha, \beta\}$, $n \in \mathbf{N}$, and it is common knowledge that if $m \in N^{(n)}$ then n is aware of the fact that $A_m = A$. Let $\bar{\Omega}$ be the set of possible probability distributions of β, so that $\bar{\Omega} = [0, 1]$, and $\nu \in \bar{\Omega}$ is the true distribution. We define by Ω_n as the set of possible samples that n can observe, $\Omega_n = \{\beta_L, \beta_H\} = \Omega$. We can think of Ω as the set of possible signals that each agent can obtain. We assume that if n is aware of m then he is also completely aware that $\Omega_m = \Omega$.

Agents make conjectures about samples of others, and then compute the frequency of each outcome from the sample. Consider a first order sample conjecture: agent n observes $\hat{\omega}_n \in \Omega$ and conjectures a set of possible samples given his observation. Let $W^{(n)}$ be this set of possible samples (consistency will in equilibrium require that $\omega_n = \hat{\omega}_n$, for each $\omega \in W^{(n)}$). Agent n also conjectures the set of sample conjectures of agent m, $W^{(n,m)}$, for $m \in \mathbf{N}$ (consistency will in equilibrium require that conjectures are only made for $m \in \mathbf{N}^{(n)}$), and so on for the higher order conjectures. Define the sample conjecture of agent n,

$$CW_n = (W^{(n)}, W^{(n,m)}, ...).$$

Let $\mathcal{U}_c = \{CN_n, CW_n \mid n \in \mathbf{N}\}$.

For each $\omega \in W^{(n)}$, each agent $n \in \mathbf{N}$ computes the maximum-likelihood estimate of ν,

$$\nu_{ML}(\omega) = \frac{\sum_{m \in \mathbf{N}^{(n)}} \mathbf{1}_{\{\omega_m = \beta_H\}}}{|\mathbf{N}^{(n)}|},$$

where $\mathbf{1}$ denotes the indicator function. It is common knowledge that agents make such estimates, so that second order sample conjectures of agent n imply maximum-likelihood estimates by other agents in his awareness, and so on at higher orders. Based on $\nu_{ML}(\omega)$, agent n rationalizes a technology. In particular, let $a^*(\nu_{ML}(\omega))$ be the optimal action given the estimated frequency of β_H, and a technology a is rationalizable for n if there exists a $\omega \in W^{(n)}$ such that $a = a^*(\nu_{ML}(\omega))$. Let S_n be the set of rationalizable technologies of agent n. The equilibrium "outcome" is now the profile $S = (S_1, S_2, S_3)$ of sets of rationalizable technologies, where S_n is a mapping from n's conjectures and sample into subsets of A. Note that if there were full awareness in this model, the notion of equilibrium would have to be weak, e.g. rationalizability. The reason is that the finite sampling doesn't allow for a stronger equilibrium notion, for example one closer to a mixed-strategy Nash equilibrium.

We now define the appropriate notion of AE which corresponds to the case where only sets S_n are verifiable *ex post* but not the actual samples.

DEFINITION **2.** The *Social Awareness Equilibrium* (SAE) is defined as follows.

SAE1 $n \in \mathbf{N}^{(n)} \subset \mathbf{N}, \forall n \in \mathbf{N}$.

SAE2N $\mathbf{N}^{(n,m)} \subset \mathbf{N}^{(n)}, \mathbf{N}^{(n,m,l)} \subset \mathbf{N}^{(n,m)}$, and so on, $\forall n, m, l \in \mathbf{N}, n \neq m$.

SAE2W (a) $W^{(n,m)}$ is a component of CW_n if and only if $m \in \mathbf{N}^{(n)}$ (similarly at higher orders),
(b) $\bar{\omega} \in W^{(n)} \Rightarrow \bar{\omega}_n = \hat{\omega}_n$,
(c) $\bar{\omega} \in W^{(n)} \iff \exists \bar{\bar{\omega}} \in W^{(n,m)}$ s.t. $\bar{\bar{\omega}}_m = \bar{\omega}_m$.

SAE3 $\omega^{(n)} \in W^{(n)} \iff a^*[\nu_{ML}(\omega^{(n)}] \in S_n^*$ and $\exists \omega' \in W^{(n)}$ s.t. $Pr_{\nu_{ML}(\omega')}(\omega) > 0$ and $Pr_{\nu_{ML}(\omega')}(\omega \mid S^*, CN_n, CW_n) > 0$, and similarly for higher order conjectures.

SAE4 There exists an $S^* = (S_1^*(\hat{\omega}_1; CN_1, CW_1), S_2^*(\hat{\omega}_2; CN_2, CW_2), S_3^*(\hat{\omega}_3; CN_3, CW_3))$ consistent with the conjectures, $\forall \hat{\omega} = (\hat{\omega}_1, \hat{\omega}_2, \hat{\omega}_3)$ with a positive probability for given ν.

The SAE is a more complicated version of AE, adapted to the present model. Note that the complication does not come from awareness alone, but rather from the interplay between awareness and the inference of samples of others from the set of their rationalizable actions. SAE1 is awareness of other agents. SAE2N is the consistency of conjectures on agents' awareness of other agents. Similarly, SAE2W imposes consistency on sample conjectures given the awareness on the agents and given the observed sample. Part (a) imposes that conjectures about others' samples be consistent with awareness and conjectures about agents. Part (b) imposes that the agent's own sample observation is consistent with his sample conjectures. Part (c) imposes the consistency of sample conjectures at higher orders.

SAE3 requires that the set of samples be consistent with inference and that it be consistent with the rationalizability. That is, suppose that an agent conjectures a sample, and infers ν_{ML} from that sample. Then such ν_{ML} also implies what other samples are possible as well. Those other samples also have to be considered as valid conjectures unless they are inconsistent with the conjectured awareness of agents and observed actions. In other words, SAE3 requires consistency of conjectured samples with the way samples are generated and with other conjectures. The rest of SAE3 is simply a definition of rationalizability. SAE4 defines the equilibrium outcome as in AE4.

The interpretation of this model is as follows. If both actions are rationalizable then there is some inefficiency in the society, but the efficient technology is being used as well. If only the efficient action is played, then the social outcome is fully efficient. If only the inefficient technology is used, then the social outcome is maximally inefficient. If each agent were alone, or only aware of himself, then this would simply be a decision problem where the agent decides what technology to use on the basis of the observed frequencies in his experiment. When agents are aware of other agents as well, then their inference is also affected by the choices of others as those choices signal the observed frequencies of their experimentation. In this setting we can study whether it can be in equilibrium that the agents misperceive the awareness of others, and the impact on equilibrium outcomes, which ties back to our story of an Italian and two Slovenes in NYC. In the first two parts of this Subsection, we study SAE. We then return to the network interpretation, and we show that there are outcomes that fully reveal the shape of G.

Trivial Awareness and Correct Awareness. We start by considering the simplest case, where no agent is aware of any other agent. This is equivalent to the case where $\mathbf{N} = \{1\}$. Note that in this case all the relevant conjectures from the definition of SAE are trivial, so we omit them from the arguments of S^*.

PROPOSITION **1.** If $\mathbf{N} = \{1\}$ then the unique SAE outcome is the one where $S_1^*(\beta_H) = \{\beta\}$ and $S_1^*(\beta_L) = \{\alpha\}$

Proof. The proof is very simple, since the "conjectured" sample now really consists of only the single draw of β observed by the agent. Thus, all the inference is made on the basis of that observation, and the characterization of equilibrium is then obvious. □

While the case of $\mathbf{N} = \{1\}$ is somewhat trivial, it is an important case as a calibration of the model. It is a case of a simple decision problem, and the agent makes precisely the same inference as in a normal statistical setup.

Next we turn to the SAE where we indeed have 3 agents, each agent is aware of all the other agents and their awareness conjectures are correct. Thus, we are speaking of the case when G is a complete social network. From now on we assume that

$$\frac{1}{2}(\beta_L + \beta_H) > \alpha \text{ and } \frac{1}{3}(2\beta_L + \beta_H) < \alpha \tag{1}$$

PROPOSITION **2.** If $\nu \in (0,1)$ there exists a unique SAE such that G is a complete social network. In this equilibrium, $S_n^*(\hat{\omega}_n) = \{\alpha, \beta\}, \forall n \in \mathbf{N}, \hat{\omega} \in \{\beta_L, \beta_H\}$. Conversely, this outcome can only be supported if G is a Complete social network.

If $\nu = 1$ then there additionally exists such SAE (in terms of G), where $S_n^*(\beta_H) = \{\beta\}, \forall n \in \mathbf{N}$.

Proof. Let $\nu \in (0,1]$. We first show that $S_n^*(\hat{\omega}_n; CW_n) = \{\alpha\}, \forall \hat{\omega}_n, \forall n$ cannot be sustained.

It is enough to consider the mapping S^* for the case when $\hat{\omega} = (\beta_H, \beta_H, \beta_H)$, and suppose $S_1^*(\beta_H) = \{\alpha\}$. We show that this is impossible.
Consider the first-order sample conjecture $\omega^{(1)} \in W^{(1)}$ that supports such an SAE. By (1) and SAE3, it has to be that $\omega^{(1)} = (\beta_H, \beta_L, \beta_L)$. Thus, $\nu_{ML}(\omega^{(1)}) = \frac{1}{3}$, which implies that also *all* the other samples consistent with $\hat{\omega}_1 = \beta_H$ should be in $W^{(1)}$. Unless, of course, they can be ruled out, i.e. $Pr_{\nu_{ML}(\omega^{(1)})}(\tilde{\omega}^{(1)} \mid CW_1) = 0, \forall \tilde{\omega}^{(1)} \neq \omega_1$. However, in order to rule out the sample $\tilde{\omega}^{(1)} = (\beta_H, \beta_H, \beta_L)$, agent 1 has to conjecture that the only possible way for 2 to play $S_2^* = \{\alpha\}$, under every higher order conjecture that justifies this play is that $\hat{\omega}_2 = \beta_L$. Let's call such conjecture by 1 $\tilde{CW}_1$. But now, clearly, one of the possible conjectures that sustain player 2 playing $S_2^* = \{\alpha\}$ is precisely the conjecture analogous to $\tilde{CW}_1$, which does not rule out $\hat{\omega}_2 = \beta_H$, which is a contradiction. For $\nu \in (0,1)$, $S_n^*(\hat{\omega}_n; CW_n) = \{\beta\}, \forall \hat{\omega}_n, \forall n$ cannot be sustained, following similar logic as before.

It is also obvious that $S_n^*(\beta_L; CW_n) \neq S_n^*(\beta_H; CW_n)$ cannot be sustained under such CN_n. The reason is that all samples have positive probability (so the SAE1-SAE4 have to hold everywhere on the sample space, conditional on the sample), so that if choices are informative of samples there always exist samples at which at least one of the choices should be in S_n^* but is not. Therefore, the only SAE candidate is the mapping $S_n^*(\hat{\omega}_n) = \{\alpha, \beta\}, \forall n \in \mathbf{N}, \hat{\omega} \in \{\beta_L, \beta_H\}$, which can in fact always be a part of an SAE. To see this observe that under such choices of technologies, one possible conjecture is that such choices are done independently of the sample observation, in which case all samples are possible, none can be ruled out, and hence both technologies have to be chosen.

For the converse, we need to check that under no other social network G, it is possible to obtain $S_n^*(\hat{\omega}_n; ...) = \{\alpha, \beta\}$. Notice that all possible G's can be reduced to finite ones. The reason is that higher-order conjectures from some order on are constant, and it is in fact irrelevant for the purpose of equilibrium outcome, at what order that happens. More precisely, if an outcome were equilibrium when a specific conjecture is made at a high order, this outcome would also be an equilibrium when that same conjecture is made at first order. Now we can already rule out many social networks.

First, every social network G where there exists an n such that $g_{n,m}^{(n)} = 0, \forall m \neq n$, can be ruled out by Proposition 1. Second, observe that every social network G where there is one agent n for whom $\exists! m \neq n$ s.t. $g_{n,m}^{(n)} = 1$, can also be ruled out. The reason behind this that such agent n has to either make a conjecture that falls in the first case, or he has to make the conjecture that $g_{m,n}^{(n,m)} = 1$ in which case he behaves as if the social network G were a complete network on $\{m, n\}$. It is immediate that (1) implies that then in SAE $S_m^*(\beta_L) = S_n^*(\beta_L) = \{\alpha\}$. Thus the only remaining social network to be considered is the *misperception* case of Proposition 3. But this case can also not be supported. The reason is that once agent n made the conjecture on awareness of the other two, i.e. $g_{k,l}^{(n,m)} = 1 \iff \{k, l\} = 1$, it is necessarily the case that the action of the other two agents will be "revealing" (albeit wrongly) of their samples. But then, S_n^* will have to consist of only one action under such G.

Finally, the mapping $S_n^*(\beta_H) = \{\beta\}, \forall n \in \mathbf{N}$ is a part of equilibrium mapping if and only if $\nu = 1$. The if part is easy to see since one sample conjecture that supports it is the correct one under which no other samples have to be considered (since $\nu_{ML}(\beta_H, \beta_H, \beta_H) = 1$). The only if part follows from above. $\square$

What is important is that when G is a complete social network, the *worse* equilibrium cannot arise. The reason can be found in the present proof: the complete social network implies that it is common knowledge that each agent always has to allow the same kind of conjectures that he is making for other agents to make. We emphasize that this moderately inefficient outcome fully reveals the social structure. Once we move away from the complete social network, such arguments no longer apply. An agent, through his wrong conjecture about others' awareness, assumes that there is asymmetry in the way agents reason. In turn, such reasoning sustains the social network. Nevertheless, the outcome of this *misperception equilibrium* again fully reveals the social structure.

Misperception We now turn to the other SAE in this model. We focus on SAE where awareness is symmetric in the sense that n is aware of m if and only if m is aware of n. We show that there exists a *misperception* equilibrium. By misperception we mean that each agent n incorrectly, albeit consistently, conjectures that the other two agents are not aware of each other but are aware only of n. We call the resulting social network G the *Social star network*: each player perceives himself to be a center of a star (or each player conjectures he is a star), and that the other two are on the periphery.[4] Clearly, a Social star network entails incorrect conjectures. At the same time, agent n incorrectly conjectures the sampling of the other two, but in conjunction with the Social star network, this is again consistent. In stark contrast with the Complete social network, the Social star network is sustained as a part of misperception equilibrium leading uniquely to the worse possible outcome.

PROPOSITION **3.** There exists an SAE such that G is a Social star network, and such that $S_n^* = \{\alpha\}, \forall n \in \mathbf{N}$. Conversely, if $S_n^* = \{\alpha\}, \forall n \in \mathbf{N}$, $\nu \in (0,1)$, and the awareness is symmetric, then G is a Social star network.

Proof. The first part of the proof goes by construction. Without loss of generality let $n = 1$ and let his agent-awareness conjectures be as in the statement of the Proposition. We have to show that even if $\hat{\omega}_1 = \beta_H$ it still can be possible that $S_1^*(\beta_H, CN_1, CW_1) = \{\alpha\}$, for all sample conjectures that can't be ruled out.

First, consider agent 2 from 1's perspective. Note that CN_1 implies that from 1's perspective agent 2 is making all his inference and sampling conjectures only by considering agents 1 and 2. Also, CN_1 implies that from 1's perspective 2 conjectures that 1 is considering only agents 1 and 2 (and this is the only conjecture consistent with $N^{(1,2)}$). Suppose that $\hat{\omega}_2 = \beta_H$. Then, by (1) and by CN_1 it has to be that $S_2^* = \{\beta\}$ and the converse holds as well. In particular, if $S_2^* = \{\alpha\}$ it has to be that $\hat{\omega}_2 = \beta_L$. Similarly for agent 3.

Thus, if $S_2^* = S_3^* = \{\alpha\}$ then it is valid that $W^{(1)} = \{(\beta_H, \beta_L, \beta_L)\}$ in which case $S_1^* = \{\alpha\}$ by (1). This construction satisfies SAE1-SAE4.

For the converse, we eliminate all other possibilities. By the same reasoning as in the proof of Proposition 2 no social network when one of the agents conjectures the awareness of others correctly is possible. On the other hand, every network where

[4]Formally, the Social star network is defined by $\mathbf{N}^{(n)} = \mathbf{N}$, and $\mathbf{N}^{(n,m)} = \{n, m\}$ for every $m \neq n$, $n, m \in \mathbf{N}$, and the same for higher order conjectures.

one of the agents is aware only of himself is also inconsistent with the prescribed outcome. This concludes the proof. □

We stress the point that the SAE of the previous proposition is an SAE even if $\nu = 1$, i.e. when technology β is in fact deterministic and superior to α.

We remark that the case of misperception is the only equilibrium case, when samples are independent, where agent n observes $\hat{\omega}_n = \beta_H$ and nonetheless plays *only* action α. One may then argue that such an agent is bound to undertake the "revealed G approach" and understand that they are playing a bad equilibrium. That argument is incorrect. First, there is no common knowledge of equilibrium here, and an equilibrium is simply a situation consistent with agents' fully rational reasoning. Even in a standard coordination game there may be common knowledge of being in a bad equilibrium and since there is no channel for credible communication rational agents can't decide to play another one - it is *an equilibrium*. Second, for n the true samples of the other agents are unobservable and n's reasoning is perfectly consistent with a situation where their samples are not independent, although the technologies work the same for all of them. That is a perfectly possible and reasonable assumption from the viewpoint of n. A constant theme of every awareness model is that the econometrician is an omniscient observer while the agents in the model are aware of the model only partially, and this is naturally also the case with the equilibrium.

3.3 SAE as a network equilibrium concept

To conclude this Section we stress one point. In this model the *social network* has a very precise meaning, since it is defined as a one-to-one mapping from fundamentals which are not reduced form expressions of some underlying non-fundamental facts. In the economic literature on networks it is imposed that agents have preferences over links, but how this is derived from true fundamentals is not explained.[5] Moreveor, the so-called "fundamentals" (links there represent friendship or communication flows or collaborative activities, etc.) are defined specific to the network model at stake. Here, the fundamentals are defined generally, and the definition of awareness as sets and conjectures is so commonly accepted that it is quite vacuous in itself. Thus, the link here is awareness, and what such awareness means is completely unambiguous and non-specific to the network interpretation. The social network is then further specified in the AE, which is an equilibrium concept that is inherited by the social network. No exogenous consideration of how networks "really work" is needed, but there are informational network effects. While SAE is weak equilibrium notion, we can make substantive predictions for a very standard social network situation. This further validates the present modeling approach.

[5]See Jackson [2005] for a survey on economics of networks.

Finally note that the social networks here are very rich objects appropriate for modeling complexity. Complexity is not only about the number of players which can be "connected" in a variety of complex ways. Complexity is also the depth of considerations that rational agents have to undertake, which is the defining feature of the present model. We return to these ideas of complexity in the last part of the next Section.

4 "Games" and Action-awareness Equilibrium

We study games where each agent may be aware only of some actions and the corresponding outcomes. The set of agents is common knowledge. The mapping from outcomes into payoffs is common knowledge as well. We focus on games without uncertainty (i.e. with complete information), so that the awareness equilibrium builds on Nash Equilibrium in the sense that agents play and conjecture best replies (given their awareness). We present this Section in terms of 2-player finite games to make notation easier to read. It is clear how to generalize the present analysis to standard games with arbitrary number of players, arbitrary actions sets, or uncertainty.

Let $\mathbf{N} = \{1, 2\}$ be the set of agents, let $\mathbf{A} = \mathbf{A}_1 \times \mathbf{A}_2$ be the set of action profiles, where $\mathbf{A}_n$ is finite for each $n \in \{1, 2\}$, $\mathbf{A}_1 = \{\underline{1}, ..., \underline{K}\}$ and $\mathbf{A}_2 = \{\bar{1}, ..., \bar{K}\}$. Thus, $a = (\underline{a}, \bar{a})$ is a typical element of $\mathbf{A}$. Denote by $\mathcal{O}$ the set of outcomes. $\mathcal{O}$ is one-to-one with $\mathbf{A}$ so that outcome $o_a \in \mathcal{O}$ corresponds to the action profile $a = (\underline{a}, \bar{a}) \in \mathbf{A}$. Utility functions of the agents are represented by a mapping $u : \mathcal{O} \to I\!R^2$, and this mapping is unambiguous. That is, an agent is aware of an outcome if and only if he is aware of the corresponding action profile and payoffs. Payoffs are denoted by $u(a) = (u_1(a), u_2(a)), \forall a \in \mathbf{A}$.

If $\mathbf{N}, \mathbf{A}, \mathcal{O}, \mathbf{u}$ were all common knowledge then this would be a standard game. In our framework, $\mathcal{U}_{ob} = \mathbf{N} \cup \mathbf{A} \cup \mathcal{O} \cup u$. Since $\mathbf{N}$ is common knowledge, the functioning of u is common knowledge, and outcomes can be identified with profiles of actions, our definition of Awareness Equilibrium can be more specific. In particular, players' awareness restricts the set of actions of each player, and players make conjectures about others' awareness and others' conjectures and so on. Denote by $A^{(n)} = A_1^{(n)} \times A_2^{(n)} \subset \mathcal{A}$ the action-awareness of player n, so that $A_1^{(n)}$ are actions of player 1 that player n is aware of, $n \in \{1, 2\}$. A first-order conjecture of agent n about m's awareness $A^{(m)}$ is denoted by $A^{(n,m)} = A_1^{(n,m)} \times A_2^{(n,m)}$, and so on. Define the conjecture of agent n by $c_n = (A^{(n)}, A^{(n,m)}, A^{(n,m,n)}, ...)$, $n \in \{1, 2\}, n \neq m$. This defines the set of conjectures,

$$\mathcal{U}_{con} = \{c_n; n = 1, 2\}.^6$$

[6]Mathematically $\mathcal{U}_{con}$ is a collection of infinite vectors, each component being a product of subsets of action sets.

We could also stick the common knowledge into $\mathcal{U}_{con}$ but this is unnecessary since those conjectures are already completely specified.

DEFINITION **3.** The *Action-awareness Equilibrium* (AAE) is defined as follows.

AAE1 $A^{(n)} \subset \mathcal{A}, n = 1, 2.$

AAE2 $A_l^{(n,m)} \subset A_l^{(n)}, A_l^{(n,m,n)} \subset A_l^{(n,m)}$, and so on, $\forall n, m, l \in \{1, 2\}, n \neq m$.

AAE3 Action a_n^* is a best reply at every component $c_{n,s}$, $s = 1, ..., \infty$, $n = 1, 2$.

AAE4 $o_{a^*} \in \mathcal{O}$, and $a^* \in \cap_{n\in\{1,2\}} \cap_{s=1,...,\infty} c_{n,s}$, where $a^* = (\underline{a}^*, \bar{a}^*)$.

Note that the above $\mathcal{U}$ does not describe a standard game because it specifically models the fact that some part of the game (i.e. actions) may not be common knowledge. Thus, $\mathcal{U}$ is truly a different and more complex object, but it contains the standard game with the appropriately specified awareness and conjectures. In order to not overcomplicate our language, we from now on refer to AAE of Γ, although what we mean is really AAE of $\mathcal{U}_{ob} = \Gamma$.

4.1 The predictive power of AAE

We illustrate our definitions with a simple example

EXAMPLE **2.** Let Γ be given by the game in Figure 1. Consider the strategy profile $(\underline{3}, \bar{3})$. Trivially, if player 1 is aware of all his actions, this cannot be an AAE because that player would deviate to $\underline{2}$. But $(\underline{3}, \bar{3})$ can be supported in AAE. For instance, take the following awareness and conjectures, $A^{(1)} = \{\underline{1}, \underline{3}, \bar{1}, \bar{2}, \bar{3}\}$, $A^{(2)} = \mathcal{A}$, $A^{(1,k,...,l)} = A^{(1)}$ for all sequences $(1, k, ..., l)$, and $A^{(2,1)} = A^{(2,k,...,l)} = \{\underline{3}, \bar{3}\}$, for all sequences $(2, k, ..., l)$ of length at least 2. It is straightforward to check that this is an AAE.

$1\backslash 2$	$\bar{1}$	$\bar{2}$	$\bar{3}$
$\underline{1}$	0, 2	3, 3	0, 2
$\underline{2}$	2, 2	2, 1	2, 1
$\underline{3}$	1, 0	4, 0	0, 1

Figure 1

When each player is aware of all $\mathcal{U}_{ob}$ and makes the correct conjecture about the other player it is as if the game were common knowledge. This implies that if we allowed for mixed strategies the existence of AAE in the present case would follow trivially from the existence of Nash equilibrium in a standard game (finite action spaces). The reason we don't consider mixed strategies is purely for the sake of simplicity, but one could clearly extend our definitions to allow for that. Nonetheless, the existence of AAE, as defined above, is not a big problem since when each agent is aware of only one action, and the conjectures about the other agent are correct, the associated outcome supports an AAE.

Every outcome can support an AAE. Without additional structure on either awareness or conjectures in the present context AAE will not have predictive power in terms of outcomes. Nonetheless, it is possible to ask the question of what collection of awareness and conjectures will justify a given outcome. This perspective may be a useful one from an econometric standpoint.

Moreover, modeling awareness is to model players with bounded cognition capabilities. Agents not being aware of all of their actions seems particularly appealing when the game being played is very large. In such a game, bounding agents' awareness seems like a very natural way of defining bounded cognition. The present model allows to parametrize how cognitively bounded agents are. This is useful not only to rule out trivial AAE, but also to provide a more specific econometric method to measure awareness from observed outcomes. We thus define the following two classes of AAE.

Definition **4.** Fix an $\ell \geq 1$. An ℓ-Action-awareness equilibrium, ℓ-AAE, is an AAE where $|A_l^{(n)}| = \ell$, $A^{(n)} = A^{(m)}$, and $c_{n,s} = c_{n,1}$, for $n = 1, 2, l = 1, 2, m \neq n$, $s = 1, ..., \infty$.

Definition **5.** Fix an $\ell \geq 1$ and let $a^* = (\underline{a}^*, \bar{a}^*)$ be a pure strategy Nash equilibrium of the standard game given by $\mathbf{N}, \mathbf{A}, \mathcal{O}, \mathbf{u}$. An ℓ^*-Action-awareness equilibrium, ℓ^*-AAE, is an ℓ-AAE where $a^* \in A^{(n)}, n = 1, 2$.

In an ℓ-AAE the cardinality of each players' awareness of own actions and every order of conjecture about the other player's action set is ℓ, $\ell \geq 1$, the conjectures are correct, and their awareness is the same. Thus, the $\ell \times \ell$ matrix given by agents' awareness is essentially common knowledge. Note that once we impose that the cardinality of every component of the conjecture is the same, it has to be in every equilibrium that all the components of each player's conjecture coincide.

We remark that by not imposing the restriction that the conjectures be correct, no more outcomes can be supported than under ℓ-AAE (where that restriction is in place). The reason is quite simple. Let's restrict, as under ℓ-AAE, to action awareness of each player that is of cardinality ℓ and let all conjectures coincide with that set. Suppose that the two players have different awareness sets and consequently different conjectures, i.e. $A^{(1)} \neq A^{(2)}$. Then, for an outcome a to be supported under AE with such structure of conjectures and awareness, it has to be that a is a Nash Equilibrium of the game defined by $A^{(1)}$ and a Nash Equilibrium of the game defined by $A^{(2)}$. Clearly, this is least restrictive when $A^{(1)}$ and $A^{(2)}$ coincide.

Before proceeding to the results, we further elaborate on the motivations behind ℓ-AAE and ℓ^*-AAE. There are two reasons as to why we consider ℓ-AAE. The first one is to provide some structure under which we can then "vary" the agents' awareness, and see how the equilibrium predictions change as agents are more or less aware. For instance, when $\ell = \underline{K} = \bar{K} = K$ we always obtain the standard game, and when $\ell = 1$ every outcome supports an equilibrium. ℓ-AAE allows us to study what happens in between. As we mentioned above, the second motivation for ℓ-AAE is to study cognitively-bounded agents. Agents are explicitly restricted to be aware of ℓ^2 action profiles, and the fact that this is common knowledge also substantially simplifies their conjectures. However, this is very different from bounded rationality where agents are aware of everything but for example only make conjectures of finite order. Under ℓ-AAE the agents are still infinitely sophisticated and consistent thinkers. That assumption can clearly be relaxed to allow for agents that are also boundedly rational, on top of being cognitively bounded.

The motivation behind ℓ^*-AAE is a bit more subtle, and has an analogy in learning. Suppose an action profile $a(\ell)$ is an ℓ^*-AAE but is not a Nash equilibrium of the underlying standard game. Assume there existed some abstract process by which agents become more aware. ℓ^*-AAE says that even if agents become aware of some Nash equilibrium profile a^* of the standard game, they could still play $a(\ell)$. Thus, Nash equilibrium profiles of the original game will be unique absorbing states of such an abstract awareness process only for ℓs for which there do not exist other ℓ^*-AAE. Another use of ℓ-AAE and ℓ^*-AAE is that these two equilibrium notions allow an econometrician or an experimentalist to measure agents' cognitive bounds.

In the following example we show that there exist non-trivial games where there is a unique Nash equilibrium in the game $\mathbf{N}, \mathbf{A}, \mathcal{O}, \mathbf{u}$, but for every $\ell < K$ every outcome can be supported not only as an ℓ-AAE but even as an ℓ^*-AAE. The case of ℓ-AAE is very simple, and the whole game can have the unique Nash equilibrium in dominant strategies.

Example **3.**

1. For each K there exists a $\Gamma = (\mathbf{N}, \mathbf{A}, \mathcal{O}, \mathbf{u}), |A_n| = K, \forall n \in \{1, 2\}$, such that Γ has a unique pure-strategy Nash Equilibrium in strictly-dominant strategies, and for each $\ell, 1 \leq \ell < K$, every outcome is sustainable as ℓ-AAE.

2. For each K there exists a $\Gamma = (\mathbf{N}, \mathbf{A}, \mathcal{O}, \mathbf{u}), |A_n| = K, \forall n \in \{1, 2\}$, such that the following holds. Γ has a unique pure-strategy Nash Equilibrium, let $(\underline{1}, \bar{1})$ be the unique NE. Then for each $\ell, 2 \leq \ell < k$, every outcome is sustainable as an l^*-AAE.

A game to prove claim 1 is quite simple. Consider the following symmetric game. Let $u_1(\underline{1}, \bar{q}) = u_2(\underline{q}, \bar{1}) = 1$, for $q = 1, ..., K$, and $u_n(\underline{p}, \bar{q}) = 0$ for all other cases. It is immediate to see that this game has a unique Nash equilibrium profile $(\underline{1}, \bar{1})$ which is a dominant-strategy equilibrium. It is also straightforward to verify that every outcome is sustainable as an equilibrium for each $l - awareness$, $1 \leq \ell < k$ is also straightforward. Figure 2 provides an example of such a game.

$_1\backslash^2$	$\bar{1}$	$\bar{2}$	$\bar{3}$
$\underline{1}$	$1,1$	$1,0$	$1,0$
$\underline{2}$	$0,1$	$0,0$	$0,0$
$\underline{3}$	$0,1$	$0,0$	$0,0$

Figure 2

The proof of claim 2 is also quite intuitive, and only slightly more complicated. Now define u as follows. Take first the matrix for the row player, $u_1(\underline{p}, \bar{q}), 1 \leq p, q \leq K$. Let $u_1(\underline{1}, \bar{q}) = 1, u_1(\underline{q}, \bar{1}) = 0, q = 1, ..., K$. For each column $p = 2, ..., K$, assign a 1 in precisely one unassigned location in such a way that the assigned 1's don't lie in only one row. This can obviously be done. Let $u_1(\underline{p}, \bar{q}) = 0$ for all the other locations. Take player 2 and do exactly the same, but also take care so that $(u_1(\underline{p}, \bar{q}), u_2(\underline{p}, \bar{q})) \neq (1, 1)$ for $(\underline{p}, \bar{q}) \neq (\underline{1}, \bar{1})$. Since the 1s assigned to columns of player 1 are not in the same row, such assignment is possible (reader can easily verify that). See Figure 3 for an example of such a game.

$_1\backslash^2$	$\bar{1}$	$\bar{2}$	$\bar{3}$
$\underline{1}$	$1,1$	$1,0$	$1,0$
$\underline{2}$	$0,1$	$1,0$	$0,1$
$\underline{3}$	$0,1$	$0,1$	$1,0$

Figure 3

Now we have to show that u has the desired properties. Clearly, the profile $s = (\underline{1}, \bar{1})$ is a pure strategy NE of the whole game. To show that this is the unique pure-strategy NE of the whole game observe that for every $(p, q) \in \{1, ..., K\}^2, (p, q) \neq (1, 1)$, at least one player gets a 0. Suppose (wlog) it is the row player 1. Then, by construction there is another column q' such that $u_1(\underline{p}, \bar{q}') = 1$ so that 1 would want to deviate.

To show that every outcome is an ℓ^*-AAE, for all $\ell < K$, observe first that if an outcome is an ℓ^*-AAE, $\ell > 2$ then it must be an $(\ell-1)^*$-AAE (reduce the supporting awareness set of each player by one action). Thus, it is enough to show the claim for $\ell = K - 1$. So take an outcome $(\underline{p}, \bar{q}) \in \{1, ..., k\}^2, (p, q) \neq (1, 1)$, and suppose that $u_1(\underline{p}, \bar{q}) = 0$ (if it is 1, then there is no deviation for player 1 anyway). This is not column 1, since there player 1 gets 1. By construction there are $K - 2$ other rows in column p such that player 1 gets 0 in those rows, and taking those $K - 2$ rows and row p also includes row 1. Similarly for player 2, so that we have constructed the awareness sets which include action 1 for both players, and no player has a profitable deviation from the profile $(\underline{p}, \bar{q})$.

Example 3 demonstrates that *many* outcomes will be supportable even as ℓ^*-AAE outcomes in non-trivial games. Note that in the above two games each player is indifferent between many outcomes, and that for a given action of the opposing player, a player is indifferent between many of his actions. This is a non-generic situation since if we perturb the payoffs of players slightly they will never be indifferent between their actions. Thus, is the "everything goes" claim from Example 3 generic or not? Our notion of genericity is very simple.

We say that a game is *generic* if the following *no-indifference* condition holds. Let Γ be a $K \times K$ game. We say that Γ satisfies the no-indifference condition if $u_1(\underline{p}, \bar{q}) \neq u_1(\underline{p}', \bar{q}), \forall p \neq p', \forall q$ (and similarly for player 2).

THEOREM **1.** Let Γ be a generic $K \times K$ game. Denote by $e_\ell(\Gamma)$ the number of distinct ℓ-AAE outcomes of Γ, for each $\ell \in \{1, ..., K\}$. Then $K^2 - 2(\ell - 1)K \leq e_\ell(\Gamma) \leq K^2 - (\ell - 1)K$, $\forall \ell \in \{1, ..., K\}$.

Proof. Let player 1 be the row player. Given an ordered set S, Denote by $S_{(r)}$ the r-th order statistic of S.
Step 1. A profile $(\underline{p}, \bar{q})$ is supportable as an ℓ-AAE if and only if

$$u_1(\underline{p}, \bar{q}) \geq \{u_1(\underline{1}, \bar{q}), ..., u_1(\underline{K}, \bar{q})\}_{(\ell)} \text{ and } u_2(\underline{p}, \bar{q}) \geq \{u_1(\underline{p}, \bar{1}), ..., u_2(\underline{p}, \bar{K})\}_{(\ell)}. \quad (2)$$

To illustrate the logic behind this result, suppose first that $\ell = 2$. Under genericity, the claim is that a strategy profile $(\underline{p}, \bar{q})$ is then sustainable as an 2-AAE if and only if

$$u_1(\underline{p}, \bar{q}) > \min_{p' \in \{1,...,K\}} u_1(\underline{p}', \bar{q}) \text{ and}$$

$$u_2(\underline{p}, \bar{q}) > \min_{q' \in \{1,...,K\}} u_2(\underline{p}, \bar{q}').$$

To see the only if part, suppose that $u_1(\underline{p}, \bar{q}) = \min_{p' \in \{1,...,K\}} u_1(\underline{p}', \bar{q})$. By genericity of Γ it is therefore $u_1(\underline{p}, \bar{q}) < u_1(\underline{p}', \bar{q}), \forall p' \neq p$. This implies that regardless of what other row p' comprises $A^{(1)}$, player 1 will at the profile $(\underline{p}, \bar{q})$ deviate to $\underline{p}'$.

To see the if part suppose that a profile $(\underline{p}, \bar{q})$ satisfies the above condition. There exist a $p' \neq p$ and a $q' \neq q$ such that $u_1(\underline{p}, \bar{q}) > u_1(\underline{p}', \bar{q})$ and $u_1(\underline{p}, \bar{q}) > u_1(\underline{p}, \bar{q}')$. Let $A^{(1)} = A^{(2)} = \{\underline{p}, \underline{p}', \bar{q}, \bar{q}'\}$, and $(\underline{p}, \bar{q})$ is an 2-AAE outcome supported by such awareness structure. Similarly, we prove the claim for general ℓ. Note that we do not need genericity in this step. End of Step 1.

By genericity of Γ, there exists a strict ordering of 1's payoffs in each column, and a strict ordering of 2's payoffs in each row.

Step 2. $e_\ell(\Gamma) \geq K^2 - 2(\ell - 1)K$.

Fix an $\ell \in \{1, ..., K\}$. By Step 1, we will minimize the number of outcomes that can be supported under ℓ-AAE by "optimally" assigning the $\ell - 1$ lowest payoffs to player 1 in each column and $\ell - 1$ lowest payoffs to player 2 in each row. An allocation which minimizes the number of outcomes supportable as ℓ-AAE is one where all these payoffs are allocated to different profiles. Since there are K columns, $\ell - 1$ worse payoffs to 1 in each column, K rows, and $\ell - 1$ worse payoffs to 2 in each row, there are in total at most $2K(\ell - 1)$ action profiles that can be eliminated. This gives the desired lower bound on $e_\ell(\Gamma)$.

Step 3. $e_\ell(\Gamma) \leq K^2 - (\ell - 1)K$.
Fix $\ell \in \{1, ..., K\}$. By Step 1, we will maximize the number of outcomes by allocating the $\ell - 1$ lowest elements of each row and each column in way which takes least space in the game matrix. That is achieved for instance by having every outcome which is the worst payoff in a given row for the column player to also be the worst payoff in the given column for the row player. Since there are K rows and columns and there are by genericity $\ell - 1$ strictly worst payoff in each, we can thus eliminate at least $K(\ell - 1)$ outcomes, which gives the desired upper bound on $e_\ell(\Gamma)$. □

Corollary 1. Let $K \geq 3$, and let Γ be a generic $K \times K$ game, then the set of ℓ-AAE outcomes of Γ is a strict superset of the set of pure-strategy Nash Equilibrium outcomes of Γ, $\forall \ell \leq \frac{K}{2}$.

Proof. A generic $K \times K$ game can have at most K pure-strategy Nash Equilibrium outcomes, and the claim follows. □

We remark that the bounds in Theorem 1 are clearly tight, which is obvious from its proof. Theorem 1 in principle says nothing about the bounds on the number of ℓ-AAE relative to the number of pure-strategy Nash Equilibria. On one hand, pure strategy Nash Equilibria may not exist. On the other hand, when pure strategy Nash Equilibria exist, it is clear that the lower bound on the number of ℓ-AAE may be in some cases improved, since every Nash Equilibrium is also an ℓ-AAE, for all ℓ. Nevertheless, the main point of Theorem 1 is that it is impossible to provide conditions which would assure that the set of ℓ-AAE is equal to the set of Nash Equilibria, if we do not correlate ℓ to the size of the game (which would hardly make sense).

In the case of ℓ^*-AAE it is to some extent obvious that the number of ℓ^*-AAE is linked to the number of pure strategy Nash Equilibria. If no pure strategy Nash Equilibrium exists, then by definition ℓ^*-AAE doesn't exist either. We show that the lower bound on the number ℓ^*-AAE is always equal to the number of Nash Equilibria, regardless of how large ℓ is. In the following theorem we denote by $floor[x]$ the largest integer that is smaller than $x \in \mathbb{R}$, and by $mod[y, r]$ the leftover from integer division of an integer y with integer r.

THEOREM **2.** Let Γ be a generic $K \times K$ game. Denote by $e_N(\Gamma)$ the number of pure strategy Nash Equilibria of Γ and by $e_{\ell^*}(\Gamma)$ the number of ℓ^*-AAE of Γ. Then $e_N(\Gamma) \leq e_{\ell^*}(\Gamma) \leq floor[\frac{K-1}{K-\ell+1}](K-\ell+1)^2 + (mod[K-1, K-\ell+1])^2 + 1$.

Proof. The lower bound is a consequence of the following simple Lemma.

LEMMA **1.** $e_N(\Gamma) = e_{\ell^*}(\Gamma)$ if and only if the following condition holds. For every profile $(\underline{p}, \bar{q})$ and every Nash Equilibrium profile $(\underline{p}^*, \bar{q}^*)$, either $u_1(\underline{p}, \bar{q}) \leq u_1(\underline{p}^*, \bar{q})$ or $u_2(\underline{p}, \bar{q}) \leq u_2(\underline{p}, \bar{q}^*)$.

Proof. The if part is obvious: regardless of what Nash Equilibrium is taken along with a strategy profile $(\underline{p}, \bar{q})$, one of the players has incentives to deviate (also by genericity) to the Nash Equilibrium strategy.

To see the only if part, take a profile $(\underline{p}, \bar{q})$ and suppose there exists a Nash Equilibrium $(\underline{p}^*, \bar{q}^*) \neq (\underline{p}, \bar{q})$ such that the above condition does not hold. Take $A^{(1)} = A^{(2)} = \{\underline{p}, \underline{p}^*, \bar{q}, \bar{q}^*\}$ and it is clear that $(\underline{p}, \bar{q})$ is an ℓ^*-AAE profile for $\ell = 2$. □

The upper bound is constructed via a "geometric" argument. Fix an $\ell, 2 \leq \ell < K$, and we show by induction on ℓ and K that $e_{\ell^*}(\Gamma) \leq floor[\frac{K-1}{K-\ell+1}](K-\ell+1)^2 + (mod[K-1, K-\ell+1])^2 + 1$. Consider first a Γ, such that $e_N(\Gamma) = 1$, and assume without loss of generality that $(\underline{1}, \bar{1})$ is the Nash Equilibrium profile.

Suppose first that $\ell = 2$. Then we can for every K do the following. By genericity of Γ all the outcomes in the row 1 and column 1 cannot be sustained as ℓ^*-AAE. Also, without loss of generality, we make a construction where as many ℓ^*-AAE profiles as possible are concentrated in the lower right hand corner of the game bi-matrix. Consider a profile $(\underline{K}, \bar{K})$. This profile can be supported if in row K there is 1 outcome which is worse for player 1, and in column K there is 1 worse outcome for player 2. Moreover, $(\underline{1}, \bar{K})$ and $(\underline{K}, \bar{1})$ have to be worse for the corresponding player (since $\{\underline{1}, \bar{1}\} \subset A^{(i)}$ by definition of ℓ^*-AAE). By the same logic, all other outcomes in K-th row and K-th column can be sustained. Similarly, in all the rows $K-1, ..., 2$ the first outcome cannot be sustained but all the others can. The same applies to the columns.

Now let $2 < \ell < K$. Exactly as before, the outcomes in rows $\ell, ..., K$ and columns $\ell, ..., K$ are sustainable. If $2\ell - K - 1 \leq 1$ then all the outcomes in rows $2, ..., \ell - 1$ and columns $2, ..., \ell - 1$ can also be sustained as ℓ^*-AAE by making them higher than $\ell - 1$ outcomes in the succeeding rows and columns. In the first row and column only the Nash Equilibrium is sustainable.

If $2\ell - K - 1 > 1$, then consider the game Γ' obtained by taking the first $\ell - 1$ rows and $\ell - 1$ columns of Γ and let $\ell' = 2\ell - K - 1$. Now, the outcomes of Γ' that are sustainable as ℓ^*-AAE of Γ must be sustainable as $(\ell')^*$-AAE of Γ', so that $e_{\ell^*}(\Gamma) \leq (K-\ell+1)^2 + e_{(\ell')^*}(\Gamma')$. The claim now follows from induction.

Note that the assumption that Γ has a unique Nash Equilibrium was made only for convenience, since if there are more Nash Equilibria, we can first re-arrange the players' actions so that all of those lie on the diagonal. $\square$

The upper bound as stated in Theorem 2 is independent of the number of Nash Equilibria. However, if a generic game has a unique Nash Equilibrium this imposes additional structure on the game, and the upper bound may never be attained. We illustrate this with the example of potential games.[7] Potential games are a very natural class to consider since a subgame of a potential game is also a potential game, and every potential game has at least one pure-strategy Nash Equilibrium. Many commonly studied games are potential games, e.g. prisoners' dilemma, congestion games, or Cournot games with quasi-linear demand. Potential games were introduced by Monderer and Shapley [1996], and we refer to that paper for more general definitions and properties.

EXAMPLE **4.** A game Γ is an *ordinal potential game* if there exists a potential function $P : \mathbf{A} \to I\!R$ which represents Γ in the following way:

$$u_1(\underline{p}, \bar{q}) - u_1(\underline{p}', \bar{q}) > 0 \iff P(\underline{p}, \bar{q}) - P(\underline{p}', \bar{q}) > 0, \text{ and}$$

$$u_2(\underline{p}, \bar{q}) - u_2(\underline{p}, \bar{q}') > 0 \iff P(\underline{p}, \bar{q}) - P(\underline{p}, \bar{q}') > 0, \forall \underline{p}, \underline{p}' \in \mathbf{A}_1, \forall \underline{q}, \underline{q}' \in \mathbf{A}_2.$$

We note that a profile $(\underline{p}, \bar{q})$ is a pure-strategy Nash Equilibrium of Γ if and only if $P(\underline{p}, \bar{q}) \geq \max\{P(\underline{p}', \bar{q}); \underline{p}' = 1, ..., K\} \cup \{P(\underline{p}, \bar{q}'); q' = 1, ..., K\}$. In particular the maximum of all elements of matrix P is a pure-strategy Nash Equilibrium of Γ.

First take the 4×4 potential game Γ with $e_N(\Gamma) = 2$, given by the following matrix P.

P	$\bar{1}$	$\bar{2}$	$\bar{3}$	$\bar{4}$
$\underline{1}$	10	0	2	6
$\underline{2}$	1	3	5	2
$\underline{3}$	2	4	6	3
$\underline{4}$	5	8	7	9

Figure 4

Let $\ell = 2$, so that by Theorem 2 the upper bound on $e_{\ell^*}(\Gamma) = 10$. Clearly, the 9 right lower corner outcomes of Γ along with the left upper corner Nash Equilibrium constitute the set of 2^*-AAE of Γ, so that the upper bound is tight in this case. Also note that it is easy to extend the example to general potential games with at least 2 Nash Equilibria and different ℓ.

[7]We thank Leeat Yariv for suggesting us to study potential games.

Consider now the potential game Γ with a unique pure-strategy Nash Equilibrium, given by the matrix $\tilde{P}$ below. The unique Nash Equilibrium of Γ is the profile $(\underline{1}, \bar{1})$.

P	$\bar{1}$	$\bar{2}$	$\bar{3}$	$\bar{4}$
$\underline{1}$	10	9	2	6
$\underline{2}$	1	3	5	2
$\underline{3}$	2	4	6	3
$\underline{4}$	9	8	7	4

Figure 5

Now observe that the path of best replies from each profile $(\underline{p}, \bar{q})$ eventually ends up in $(\underline{1}, \bar{1})$. This is in fact a property of potential games with a unique Nash Equilibrium. At some point such path enters either column or row 1, suppose that the path enters column 1 in row p (in $\tilde{P}$, $p = 3$). But this implies that $P(\underline{p}, \bar{1}) > P(\underline{p}, \bar{q}), q > 1$, so that no element in row p can be sustainable as an ℓ^*-AAE outcome, which means that the upper bound may never be attained in a potential game with a unique Nash Equilibrium. Nonetheless, the additional structure imposed by uniqueness of Nash Equilibrium eliminates only one additional row (or column), so that in a large game this effect is negligible.

Just to show that uniqueness of Nash Equilibrium doesn't necessarily have such effect in a general game, we provide the following 4×4 game with a unique Nash Equilibrium and where the upper bound from Theorem 2 obtains. Note that this requires cycles in best replies to pure strategies which can not occur in a potential game.

[GAME: 4×4, unique Nash Equilibrium and upper bound]

Before concluding this Section, we discuss the concept of ℓ^*-AAE a bit further. Namely, one could motivate other notions similar to ℓ^*-AAE, but where for instance a whole set of Nash Equilibrium profiles is in the players' awareness set. As defined now, ℓ^*-AAE is the weakest such equilibrium, so that the upper bound here is also an upper bound on the number of outcomes supportable by stronger versions. On the other hand, it is clear that the lower bound is unaltered and is still tight under stronger versions of ℓ^*-AAE.

4.2 AE as a critique of the literature on *Unawareness*

We now discuss the relationship of our model to the existing literature on epistemic foundations of *unawareness*, which focuses entirely on games and awareness of actions, so that AAE is the appropriate notion for comparison.[8] The main point of

[8] See Feinberg [2005] and the references therein, in particular, Heifetz et al [2004], Li [2004], Modica and Rustichini [1999], Dekkel et al. [1998]

this subsection is a critique of the focus of that literature. That literature focuses on axiomatic constructions of awareness and conjectures. In examples, specific constructions are always *assumed* on agents, and the *implications* of such constructions are analyzed. There are two important points. First, assuming specific awareness on the agents drives the results, while awareness is an entirely unobservable attribute. Second, contrary to our approach, implications do not need to be consistent with imposed awareness constructions and agents' reasoning. That is, allowable awareness structures are not affected by the outcome. On the one hand, in all this literature the awareness constructions satisfy some weakening of common knowledge. On the other hand, the l-AAE is consistent with common knowledge of the smaller game (for example, the trivial game with one action), so that l-AAE satisfies all the axioms provided by the literature. Therefore, it is clear from our analysis that for each specific set of axioms proposed by the literature, if all constructions are allowed, everything is possible. Thus, we believe, that either that literature should focus on specific examples where it is clear from empirical evidence why certain awareness constructions and axioms should be relevant. Or, alternatively, the literature should provide some meaningful general restrictions on awareness constructions that reduce the set of possible outcomes. Our approach is congruent with both views.

To illustrate our critique, we focus on Feinberg [2005] which is the most related to our paper, and in terms of applying unawareness to specific class of problems it also subsumes the previous literature. As said, under the weak set of axioms which we discuss below, Feinberg [2005] allows all awareness and conjecture constructions satisfying internal consistency but not internal consistency with actions or with the outcome.[9] In other words, there is no link between these constructions and the observables so that effectively no constructions can be ruled out.

Specifically, these axioms provided by Feinberg [2005] are the Weak Inter-awareness Axiom (WI), the Positive Introspection of Awareness Axiom (PI) and the Self Awareness Axiom (SA). Our AE1 is equivalent to WI, and AE2 is equivalent to PI. SA requires that each agent is aware of his own existence and that being aware of something at some higher order awareness implies being aware of himself. As demonstrated by Example 1 we do not need to require SA because self awareness is implied by our equilibrium requirement AE4 that awareness must be consistent with the observed outcome. Importantly AE4 also rules out many awareness constructions that are inconsistent with the observables.

the epistemic foundation of unawareness literature provides axioms which correspond to interpretations of the notion of awareness as the ability to reason about subsets of $\mathcal{U}$. However, awareness and conjecture structures satisfying WI, PI and

[9] The strong set of axioms he provides is on one hand *ad-hoc* and on the other hand it is still a weakening of common knowledge so that the predictions under those will be a strict superset of predictions of l-AAE.

SA are in our case eliminated if they are inconsistent with the outcome. Consistency with the outcome eliminates many (indeed most) of them.

To operationalize his notion of *unawareness* Feinberg [2005] defines a game with incomplete awareness as a normal form game and an awareness construction which satisfies WI,PI and SA. The awareness construction defines for each player an infinite collection of games, one game at each order of conjecture of each player, describing how at that order of conjectures the player perceives the game the opponents are playing. Each player conjectures an infinite collection of profiles of mixed strategies, one profile for each game, and those profiles need not coincide. Such a collection of profiles of mixed strategies is said to be an *Extended Nash Equilibrium* if at every order of conjecture the strategies assigned are best replies. The interpretation is that each player, at each order conjectures, is playing a best response in the corresponding game. The outcome that "actually happens" in an Extended Nash Equilibrium doesn't need to be consistent with agents' higher-order conjectures. This reasoning is completely different from our approach where the conjectures on the profile of strategies are by definition correct in equilibrium, and hence constant across all orders of conjectures. Nonetheless, every l-AAE is an Extended Nash Equilibrium, and there exist awareness constructions under which the outcome of Extended Nash Equilibrium is not supportable as an l-AAE, for $l > 1$. The predictive power of the Extended Nash Equilibrium is thus weaker than that of l-AAE, and we have already shown that the predictive power of l-AAE is so weak that it requires careful justification of high levels of cognition (high l) in order to even narrow the set of l-AAE closer to the set of Nash Equilibria.

Example **5.** To see the difference between AAE and the Extended Nash of a game consider the game from Figure 1, which is in fact taken from Feinberg [2005].

Extended Nash:$(\underline{3}, \bar{3})$ is an extended Nash equilibrium of the the game of incomplete awareness defined by the normal form game in Figure 1 and an awareness structure where 1 and 2 are both aware of all the actions (see Feinberg [2005]).

AAE: An AAE equilibrium of the game in Figure 1 where players 1 and 2 play $(\underline{3}, \bar{3})$ exists if player 1 is aware at most of $\{\underline{1}, \underline{3}, \bar{1}, \bar{2}, \bar{3}\}$, as we pointed out in Example 2. Regardless of what higher order conjectures the players make, $(\underline{3}, \bar{3})$ is not an AAE if each player is aware of all the actions (and hence also not an l-AAE).

4.3 A discussion of evolutionary perspective on cognition

[This subsection is very preliminary, and will be substantiated. Also, we are *aware* of existing work (e.g. by Binmore and by Samuelson, by Rubinstein, by Young, and others) which we should carefully review and cite.]

Here we provide a brief sketch of how the AAE may be a useful modeling tool. The discussion in this subsection is intended to motivate that AAE is not merely a device to point out the deficiencies of existing literature. But to develop these examples in depth is beyond the scope of the present paper. It is our intention to address the questions outlined here in our future work.

Consider for example two species, $\mathcal{S}_1$ and $\mathcal{S}_2$, that compete in some evolutionary game Γ_E (to be specified in particular examples). Nature is an optimizing *designer*, allocating finite intellectual resources between different aspects of intelligence of each species. That is, think of each species having some budget of intellectual units that have to be allocated between perceptive ability (cognition), memory and computation, and depth of strategic reasoning. One could in principle assume non-linear "production functions" of these processes, so that higher levels of say perceptive ability are more and more costly at the margin. Nature optimizes in view of the evolutionary game Γ_E played by the two species. Since there are no beliefs, one possible specification of utility maximization by members of species is a min max utility over the set of possible outcomes that can occur in Γ_E. As discussed above, l-AAE gives a pessimistic view of the extent of the set of outcomes that are possible under cognition l (in the sense of it being largest possible for that level of cognition). Thus it is precisely the right notion consistent with the min max criterion. Of course, to make such approach operational, we need to add bounded rationality to l-AAE (quantifying memory, computation power, and depth of conjectures). Incorporating such quantifiable aspects of bounded rationality into our l-AAE will be no harder than in the standard model. Also dynamics will play a role in such model, but the idea is to think of stochastic equilibria where the state of the world is characterized by the $l \times l$ submatrix of Γ_E that the species are aware of (it is the fundament of this approach that l is constant, as well as allocations into other aspects of intelligence). Our AE thus *naturally* relates to ideas of complexity and evolution.

In the present paper we limit our discussion to the following simple example.

[Here comes a very simple example]

5 Conclusion

We proposed a simple definition of awareness as a notion of equilibrium. This notion can be used to show meaningful implications of awareness on economic outcomes. Our approach demonstrates a critique of epistemic approaches to unawareness. Our notion of Awareness Equilibrium is simple, so that incorporating bounded rationality is no more difficult than in standard models. We propose a research agenda in which the level of awareness is determined in an evolutionary game, and our definition of Awareness Equilibrium is a building block of the model.

Our approach is also congruent with providing much more specific and empirically motivated models of awareness for particular economic situations. But in order to do that one has to turn to empirical or experimental data. In the absence of such data, awareness must in equilibrium be a simple notion which says that within the world that each agent perceives, his perceptions are not proven wrong by the outcome.

Appendix A.

Recall that $\mathcal{O}$ is the set of outcomes, agents in $\mathbf{N}$ have preferences over outcomes of which they are aware. In a setting without uncertainty each agent is aware of preferences of every agent he is aware of on the subset of outcomes he is aware of. There is a mapping O from actions $a \in \mathbf{A} = \mathbf{A}_1 \times \mathbf{A}_2 \times ... \times \mathbf{A}_N$ into $\mathcal{O}$, agents are aware of the mapping on the subset of actions they are aware of. An r-th order conjecture of agent $n \in \mathbf{N}$ is a mapping

$$c_{n,r} : C_{n,r} \to \{0, 1\},$$

where $C_{n,r} \subset \bar{\mathbf{N}}^r \times \mathcal{U}$ and $r \geq 1$. For example $c_{n,2}^{(k,l)}$ is a conjecture that n makes about k's conjecture of l's awareness. Now we can formally state the definition of AE.

AE1 $\emptyset \neq \mathcal{U}_n \subset \mathcal{U}$, $n \in \mathbf{N}$ and $C_{n,r} = (\mathbf{N} \cap \mathcal{U}_n)^r \times \mathcal{U}_n$, and this is common knowledge.

AE2 $\left[c_{n,r}^{(k_1,\dots,k_r)}\right]^{-1}(1) \subset \left[c_{n,r-1}^{(k_1,\dots,k_{r-1})}\right]^{-1}(1)$, the two sets are equal if $k_r \in \{k_{r-1}, 0\}$, and
$k_r \in \left[c_{n,r-1}^{(k_1,\dots,k_{r-1})}\right]^{-1}(1)$, and this is common knowledge.

AE3 Each $n \in \mathbf{N}$ chooses his action a_n optimally from $\mathbf{A}_n \cap \mathcal{U}_n$, and that the choice of action is made in this way is also conjectured by agent n at every order of his conjectures.

AE4 There exists an outcome $o = O\left(a_1, ..., a_N \mid \left[c_{n,r}^{(k_1,\dots,k_r)}\right]^{-1}(1)\right)$ $\forall n \in \mathbf{N}$, $\forall r \geq 1$.

References

[1] Battigalli P. and P. Guaitoli[1988]: "Conjectural Equilibrium", mimeo.

[2] Dekel B., B. Lipman, and A. Rustichini[1998]: "Standard state-space models preclude unawareness", *Econometrica*, **66**, 159-173.

[3] Ellison G. and D. Fudenberg[1995]: "Word-of-Mouth Communication and Social Learning", *The Quarterly Journal of Economics*, **110**,93-125 .

[4] Fagin R. and J. Y. Halpern[1988]: "Belief, awareness and limited reasoning", *Artificial Intelligence*, **34**, 39-76.

[5] Feinberg Y.[2005]: "Games with Incomplete Awareness", mimeo. Stanford Research Paper 1894.

[6] Fudenberg D. and D. K. Levine[1993]: "Self-Confirming Equilibrium", *Econometrica*, **61**, 523-545.

[7] Halpern J. Y.[2001]: "Alternative semantics for unawareness", *Games and Economic Behavior*, **37**, 321-339.

[8] Heifetz A., M. Meier, and B. Schipper[2004]: "Interactive Unawareness", Journal of Economic Theory, forthcoming.

[9] Jackson M. O.[2005]: "The Economics of Social Networks," Proceedings of the 9th World Congress of the Econometric Society, ed. R. Blundell, W. Newey, and T. Persson, Cambridge University Press, fortchoming.

[10] Kawamura E.[2004]: "Competitive Equilibrium with Unawareness in Economies with Production", *Journal of Economic Theory*, forthcoming.

[11] Li Jing[2003]: "A generalized correspondence model with unawareness", mimeo., University of Wisconsin-Madison.

[12] Modica S. and A Rustichini[1999]: "Unawareness and partitional information structures", *Games and Economic Behavior*, **27**, 265-298.

[13] Rubinstein A. and A. Wolinsky[1990]: "Rationalizable Conjectural Equilibrium Between Nash and Rationalizability", mimeo.