Disagreement Under Almost Common Knowledge of Rationality*

Emiliano Catonini†

Tatiana Mayskaya‡

February 2, 2020

Abstract

Two agents sincerely exchange their best guesses about the state of the world infinitely many times. When each agent places a small positive probability on the event that her opponent is of some finite level of reasoning and initial disagreement is large enough (that is, private signals are strong and different), permanent and large disagreement is possible even for infinitely sophisticated agents.

Keywords: disagreement, almost common knowledge, level-k reasoning

JEL Codes: D83

*We are grateful to Marciano Siniscalchi and Antonio Penta for helpful comments. We also received useful suggestions from audience at the Stony Brook Game Theory conference, the Lisbon Meeting in Game Theory and Applications, the Lancaster Game Theory conference, and SAET in Ischia. The study has been funded by the Russian Academic Excellence Project '5-100'.

†International College of Economics and Finance, National Research University Higher School of Economics. Postal address: 11 Pokrovsky Boulevard, 109028 Moscow, Russia. Email: emiliano.catonini@gmail.com

‡International College of Economics and Finance / Faculty of Economic Sciences, National Research University Higher School of Economics. Postal address: 11 Pokrovsky Boulevard, 109028 Moscow, Russia. Email: tmayskaya@gmail.com

1 Introduction

Aumann (1976)'s impossibility to disagree result has always been a source of uneasiness among economists. Intuitively, if two agents start a debate with very different opinions on a matter, it is unlikely that they will eventually reach an agreement. However, Aumann (1976) proved that if two agents share the same prior and have common knowledge of their information partitions, their rationality (that is, both agents use Bayes' rule to update their beliefs) and their posteriors, then these posteriors must be the same (they cannot agree to disagree).

We do not question neither the assumption of common prior[1], nor the assumption that agents' information partitions are common knowledge. Instead, we assume the agents have some doubts about *each other*'s higher-order rationality. Loosely speaking, we show that even if these doubts are very small, the agents might severely disagree.

Imagine two agents arguing about some hypothesis. Suppose agent 1 received a lot of evidence in favor of the hypothesis, while agent 2 received a lot of evidence against it. At each moment they simultaneously announce their best guess whether the hypothesis is true or false. We assume the announcements are sincere, that is, if at the current moment an agent believes it is more likely that the hypothesis is true, she announces that it is true, and vice versa. As was shown in Cave (1983), if the agents' rationality is common knowledge, these agents eventually will agree on which state is more likely. That behavior relies on the assumption of *common knowledge of rationality*: agent 1 knows that agent 2 uses Bayes rule, agent 1 knows that agent 2 knows that agent 1 uses Bayes rule, agent 1 knows that agent 2 knows that agent 1 knows that agent 2 uses Bayes rule, etc. Placing a *small* doubt on a *high* order of rationality for the opponent, we show that even infinitely sophisticated agents might never come close to an agreement.[2] More precisely, we fix three things — the size of the doubt, the order of rationality this doubt is placed upon, the distance d between posteriors that characterizes the size of ex post disagreement — and we find the following: If initial disagreement is large enough, the disagreement remains perpetually larger than d. The key intuition why the size of ex post disagreement is not bounded above by the size of the doubt is that the doubt actually grows larger and larger as time passes without ever reaching an agreement. The doubt grows because an agent with strong evidence in favor of her view finds it implausible that the other has equally strong evidence in

[1] Morris (1994) explains why heterogeneous prior might lead to disagreement.

[2] With a ring game, Kneeland (2015) estimated that 71% of her subjects have level 2 reasoning or higher, while only 22% of her subjects have level 4 reasoning or higher. With a beauty contest game, Bosch-Domenech et al. (2002) came to similar conclusions, adding that subjects are either of lower levels (1,2,3) or infinity.

favor of the opposite view, and then the only way to rationalize disagreement is to lower expectations regarding the cognitive ability of the opponent.[3]

Our result is in the spirit of Rubinstein's Email Game (Rubinstein (1989)) where the prediction under "almost common knowledge" of an event, — rationality in our case and an external state in the Email Game — is very different from the prediction that assumes common knowledge of that event. In our case, "very different" means not only that the agents with sufficiently strong initial disagreement disagree forever but that their posteriors remain arbitrarily far apart. Similar to Rubinstein, we show that the difference in the predictions remains significant no matter how close we get to common knowledge of the event.

Our paper belongs to the stream of literature that investigates the robustness of Aumann (1976)'s impossibility to disagree result. Geanakoplos and Polemarchakis (1982) showed that common knowledge of posteriors can be achieved by communicating the posteriors back and forth. Cave (1983) and Bacharach (1985) proved the impossibility to disagree also when agents take actions of arbitrary nature, not necessarily communicating their posterior beliefs. We start where they finished: their conclusions guarantee agreement under common knowledge of rationality in our benchmark model where the agents take binary actions by communicating their guesses about a binary state of the world.

Monderer and Samet (1989) generalize Aumann (1976) by relaxing the assumption of common knowledge of posteriors: if agents' posteriors are common p-belief, then these posteriors can differ by at most $2(1-p)$. In our model the agents do not achieve "almost common knowledge" of posteriors, so their posteriors can be significantly different.

Following Rubinstein (1989), Weinstein and Yildiz (2007) and many others, we study the impact of higher-order beliefs on the final outcome. Similar to Strzalecki (2014), we adopt level-k thinking to close the gap between common sense intuition and game-theoretical analysis. Strzalecki (2014) explains coordination in Rubinstein's Email Game, we explain disagreement in a communication model.

Our disagreement result in a Bayesian learning model contrasts with a series of consensus results in naive learning models (DeMarzo et al. (2003), Golub and Jackson (2010) and others). In these models a learning rule treats the opponent's opinion in a similar way each period, eventually making the agents converge. In contrast, our Bayesian agents discount the opponent's opinion over time because of the growing doubt regarding the cognitive ability of the opponent.

[3]We elaborate on the difference between attainable disagreement with and without level updating at the end of Section 3 — see in particular Figure 1.

2 Setup

The state of the world is binary, we denote it as $\theta \in \{0,1\}$. There are two agents, 0 and 1, who share a common prior belief about θ. Time is discrete and infinite: $t = 0,1,2,\ldots$. At $t = 0$, agent $i \in \{0,1\}$ privately observes the realization of a signal s^i distributed according to a cumulative distribution function F_i^θ. Signals s^0 and s^1 are independent conditionally on the state of the world. The distributions F_i^θ, $i \in \{0,1\}$, $\theta \in \{0,1\}$, are common knowledge. At each $t \geq 1$, the agents publicly and simultaneously announce their guesses about the state of the world: 0 or 1. We assume that, in each period, the agents have the incentive to guess the state *sincerely*.[4] We abstract away from any strategic considerations whereby an agent lies to elicit information from her opponent.

Now we propose a model of how the agents update their beliefs. This model departs as little as possible from common knowledge of rationality. We say that an agent is rational when, in each period, she updates her belief about the state and the level of reasoning of the opponent using Bayes rule and all the available information. We focus on rational agents who believe, with probability $1-\varepsilon$, that the opponent is rational, believes that the other is rational, and so on; and with probability ε, that the opponent is a "level k reasoner", for some $k \geq 1$. The definition of a level k agent is inductive. A level 1 agent updates the prior with her private signal and ignores the opponent's guesses; a level 2 agent updates the prior with her private signal and interprets the opponent's guesses as coming from a level 1 agent; and so on. A level ∞ agent is a rational agent for whom rationality is common knowledge. Therefore, our agents believe the opponent is of level ∞ with probability $1-\varepsilon$ and of level k with probability ε.

To complete the description, one should add assumptions on how different types of agents revise their beliefs after observing a zero probability event. However, we do not need such assumptions because in our analysis the agents will never observe a zero probability event along the permanent disagreement path (disagreement never comes as a complete surprise under the signal structure we use for our existence result).

The level k of reasoning can have two interpretations. Level 1 can be seen as a rational agent who does not believe in the rationality of the opponent and therefore disregards the opponent's guesses as noise. Alternatively, closer to the spirit of level-k reasoning proposed by Stahl and Wilson (1995), level 1 can be seen as an agent whose cognitive ability does not allow her to process the informational content of the

[4] The guessing strategy of an agent who believes that the two states are equally likely is immaterial for our analysis.

opponent's guesses. A similar distinction can be done at all levels.[5] A level 1 agent can also be viewed as a "committed" type who, for some reason, never changes her action.

The cognitive hierarchical (CH) theory proposed in Camerer et al. (2004) adopts a distribution f over levels of reasoning that represents both the realized distribution in the population and, truncated at level $k-1$, the belief of a level k agent over the level of her opponent. (If the distribution is concentrated above $k-1$, the truncation is substituted by probability 1 on level $k-1$.) We depart from the CH model in two ways. First, we do not require the probability distribution f over the levels of thinking to be consistent with the realized distribution of levels of thinking in the population. Instead, we assume that everybody is of "level ∞". Alternatively, one might prefer to interpret this assumption as if f **is** consistent with the realized distribution but we focus only on a subsample of ∞-level agents from the population. Focusing only on such agents is natural for our goal: if the "smartest" people cannot agree, we should not be surprised to observe disagreement among the others. Second, in the CH model the level ∞ agents assign probability to other level ∞ agents whose beliefs are represented by the f distribution as well. This induces "higher-order doubts" of the following kind: "I think that my opponent thinks that I might not be level infinite..." Our disagreement result does not require to introduce this sort of reasoning; therefore, we picked the closer-to-common-knowledge hypothesis that the opponent is not viewed as having this kind of doubts.

We end this section with an important observation to which we come back again after we prove the disagreement result. Our proof is invariant to how we define our agents' beliefs about rationality, as long as they put probability $\varepsilon > 0$ to level k reasoner. In particular, we could have picked the CH model with $f(1) = \ldots = f(k-1) = 0$ and taken level ∞ agents (in the sense of this CH model) as *our* agents. Or we could have assumed common p-belief of rationality (Hu (2007)) for large p (which entails positive, decreasing probability on all levels).

3 Disagreement Result

We want to find sufficient conditions under which agent 0 guesses "zero" all the time, agent 1 guesses "one" all the time, and the distance between their posteriors remains large. This means that agent 0 receives strong evidence in favor of state 0, while agent 1 receives strong evidence in favor of state 1. In the benchmark model

[5] Friedenberg et al. (2018) distinguishes between "rationality bound" and "strategic bound", with possible positive gap between the latter and the former. By assuming sincere guessing, we can abstract away from this gap.

when rationality of the agents is common knowledge (that is, when $\varepsilon = 0$), permanent disagreement is possible only in a symmetric situation where the agents' posteriors both converge to 0.5. In our model, the belief updating process of an agent is interrupted by the *growing doubt* that her opponent is not of level ∞.

Theorem 1. *There exists a signal structure such that for any non-degenerate prior about θ, any level $k \geq 1$, any positive amount of doubt $\varepsilon \in (0,1)$, and any final level of disagreement $d \in (0,1)$, we can always find $\delta > 0$ so that δ-extremists permanently disagree with each other, and the distance between their posteriors exceeds d forever. Our agents are called δ-extremists if agent 0's posterior satisfies* $\Pr(\theta = 1 \mid s^0) < \delta$ *and agent 1's posterior satisfies* $\Pr(\theta = 1 \mid s^1) > 1 - \delta$.

The idea behind the proof is the following. As time goes and disagreement persists, each agent fails to rule out the possibility that her opponent is of level k and received relatively *mild* evidence in favor of the opposite state. On the other hand, the agent does realize that, if the opponent is of level ∞, her insistence can only be justified by strong evidence in favor of the opposite state. But when the agent herself has strong evidence in favor of her position, she assigns much higher probability to the first event compared to the second.

$$\begin{aligned} &s^1 \text{ is strong in favor of } \theta = 1 \\ &\qquad \Rightarrow \quad \Pr\left(s^0 \text{ is strong in favor of } \theta = 0\right) \ll \Pr\left(\begin{smallmatrix} s^0 \text{ is weak in favor of } \theta=0, \\ \text{agent 0 is of level } k \end{smallmatrix}\right) \end{aligned}$$

That belief fuels permanent disagreement.

Proof of Theorem 1. The signal distributions F_i^θ are defined in the following way:

$$\Pr(s^i < x \mid \theta = 1) = x^2, \quad \Pr(s^i < x \mid \theta = 0) = 2x - x^2, \quad x \in [0,1]. \tag{1}$$

Denote by $h_t = (0,1)^t$ a length-t history of disagreement when agent 0 guesses "zero" and agent 1 guesses "one" in all periods $1, \ldots, t$. Without loss of generality, we focus on agent 1. Denote by L_0 the level of reasoning of agent 0. Let us fix any

$0 < \Delta_1 < \Delta_2 < 1$ (we will impose additional restrictions on them later).

$$\begin{aligned}
\Pr(\theta = 1 \mid h_t, s^1) \overset{(1^*)}{=} \; & \Pr(\theta = 1 \mid h_t, s^1, s^0 < \Delta_1)\Pr(s^0 < \Delta_1 \mid h_t, s^1) \\
& + \Pr(\theta = 1 \mid h_t, s^1, \Delta_1 \le s^0 < \Delta_2)\Pr(\Delta_1 \le s^0 < \Delta_2 \mid h_t, s^1) \\
& + \Pr(\theta = 1 \mid h_t, s^1, s^0 \ge \Delta_2)\Pr(s^0 \ge \Delta_2 \mid h_t, s^1) \\
\overset{(2^*)}{\ge} \; & \Pr(\theta = 1 \mid s^1, s^0 = 0)\Pr(s^0 < \Delta_1 \mid h_t, s^1) \\
& + \Pr(\theta = 1 \mid s^1, s^0 = \Delta_1)\Pr(\Delta_1 \le s^0 < \Delta_2 \mid h_t, s^1) \\
& + \Pr(\theta = 1 \mid s^1, s^0 = \Delta_2)\Pr(s^0 \ge \Delta_2 \mid h_t, s^1) \\
\overset{(3^*)}{\ge} \; & \Pr(\theta = 1 \mid s^1, s^0 = \Delta_1)\frac{\Pr(\Delta_1 \le s^0 < \Delta_2 \mid h_t, s^1)}{\Pr(s^0 < \Delta_1 \mid h_t, s^1) + \Pr(\Delta_1 \le s^0 < \Delta_2 \mid h_t, s^1)} \\
\overset{(4^*)}{\ge} \; & \Pr(\theta = 1 \mid s^1, s^0 = \Delta_1)\frac{\Pr(\Delta_1 \le s^0 < \Delta_2, L_0 = k \mid h_t, s^1)}{\Pr(s^0 < \Delta_1 \mid h_t, s^1) + \Pr(\Delta_1 \le s^0 < \Delta_2, L_0 = k \mid h_t, s^1)} \\
\overset{(5^*)}{\ge} \; & \Pr(\theta = 1 \mid s^1, s^0 = \Delta_1)\frac{\Pr(\Delta_1 \le s^0 < \Delta_2, L_0 = k \mid s^1)}{\Pr(s^0 < \Delta_1 \mid s^1) + \Pr(\Delta_1 \le s^0 < \Delta_2, L_0 = k \mid s^1)} \\
\overset{(6^*)}{>} \; & \frac{1+d}{2}
\end{aligned}$$

Equality (1^*) holds by the law of total probability. To show inequality (2^*), we need one simple fact:

Fact 1. *Let X be a random variable continuously distributed on the support $[a,b]$. If function $g(x)$ is increasing in x, then $\mathbb{E}[g(X)] > \mathbb{E}[g(X) \mid X \le x]$ for any $x \in [a,b)$.*

Then inequality (2^*) becomes evident once for any $0 \le a < b \le 1$ we show

$$\Pr(\theta = 1 \mid h_t, s^1, a \le s^0 < b) = \int_a^b \Pr(\theta = 1 \mid h_t, s^1, s^0)\Pr(s^0 \mid h_t, s^1, a \le s^0 < b)\, ds_0$$

$$= \int_a^b \Pr(\theta = 1 \mid s^1, s^0)\Pr(s^0 \mid h_t, s^1, a \le s^0 < b)\, ds_0 \overset{(i)}{\ge} \Pr(\theta = 1 \mid s^1, s^0 = a),$$

where inequality (i) follows from Fact 1 and the observation that conditional probabilities $\Pr(\theta = 1 \mid s^1, s^0)$ are increasing in s^0.

Since conditional probabilities $\Pr(\theta = 1 \mid s^1, s^0)$ are increasing in s^0, inequality (3^*) follows from Fact 2, the discrete version of Fact 1:

Fact 2. *If* $x_1 \le \ldots \le x_I$, $\sum_{i=1}^{I} \omega_i = 1$, $\omega_i \ge 0$, *then* $\sum_{i=1}^{I} \omega_i x_i \ge \frac{\omega_1 x_1 + \omega_2 x_2}{\omega_1 + \omega_2}$.

Inequality (4*) holds because

$$\Pr(\Delta_1 \le s^0 < \Delta_2,\ L_0 = k \mid h_t,\ s^1) \le \Pr(\Delta_1 \le s^0 < \Delta_2 \mid h_t,\ s^1).$$

Intuitively, inequality (5*) follows from the observation that agent 0 of level k reports "zero" all the time if she receives $\Delta_1 \le s^0 < \Delta_2$ (assuming Δ_2 is low enough, see assumption (2)), despite agent 1 reporting "one". Then history h_t cannot decrease the relative probability of $(\Delta_1 \le s^0 < \Delta_2,\ L_0 = k)$.

To prove (5*) formally we need Claim 1 which describes the behavior of a level k agent.

Claim 1. *For any* $k \ge 1$, *there exist* $x_k^0 \in (0,1)$ *and* $x_k^1 \in (0,1)$ *such that agent* $i = 0$ *(*$i = 1$*) of level* k *with signal* $s^0 \le x_k^0$ *(*$s^1 \ge x_k^1$*) considers state* i *more likely at every history* h_t *(*$t \ge 0$*).*

We postpone the proof for Claim 1 to the end.

Assuming

$$\Delta_2 = x_k^0, \tag{2}$$

we get (5*) from

$$\frac{\Pr(\Delta_1 < s^0 < \Delta_2,\ L_0 = k \mid h_t,\ s^1)}{\Pr(s^0 < \Delta_1 \mid h_t,\ s^1)} = \frac{\Pr(h_t \mid \Delta_1 < s^0 < \Delta_2,\ s^1,\ L_0 = k)\Pr(\Delta_1 < s^0 < \Delta_2,\ L_0 = k \mid s^1)}{\Pr(s^0 < \Delta_1,\ h_t \mid s^1)}$$
$$\overset{\text{Claim 1}}{=} \frac{\Pr(\Delta_1 < s^0 < \Delta_2,\ L_0 = k \mid s^1)}{\Pr(s^0 < \Delta_1,\ h_t \mid s^1)} \ge \frac{\Pr(\Delta_1 < s^0 < \Delta_2,\ L_0 = k \mid s^1)}{\Pr(s^0 < \Delta_1 \mid s^1)},$$

where by $\Pr(h_t \mid \Delta_1 < s^0 < \Delta_2,\ s^1,\ L_0 = k)$ we mean the probability that level k agent 0 reports "zero" in all periods $1, \ldots, t$ after observing "one" in all periods from agent 1.

Informally, we get (6*) in two steps. First, we choose Δ_1 sufficiently close to 0, so that

$$\frac{\Pr(\Delta_1 \le s^0 < \Delta_2,\ L_0 = k \mid s^1)}{\Pr(s^0 < \Delta_1 \mid s^1) + \Pr(\Delta_1 \le s^0 < \Delta_2,\ L_0 = k \mid s^1)}$$

is close to 1. Then assuming s^1 is such that

$$\Pr(\theta = 1 \mid s^1) > 1 - \delta \tag{3}$$

and taking δ close to 0, we make sure $\Pr(\theta = 1 \mid s^1,\ s^0 = \Delta_1)$ is close to 1 as well.

Formally,

$$\Pr(\theta = 1 \mid s^1,\, s^0 = \Delta_1)\frac{\Pr(\Delta_1 \le s^0 < \Delta_2,\ L_0 = k \mid s^1)}{\Pr(s^0 < \Delta_1 \mid s^1) + \Pr(\Delta_1 \le s^0 < \Delta_2,\ L_0 = k \mid s^1)}$$

$$\stackrel{(*)}{=} \frac{\frac{\Pr(\theta=1)s^1\Delta_1}{\Pr(\theta=1)s^1\Delta_1+\Pr(\theta=0)(1-s^1)(1-\Delta_1)}}{\frac{\Pr(\theta=1)s^1\Delta_1^2+\Pr(\theta=0)(1-s^1)\Delta_1(2-\Delta_1)}{\varepsilon\left(\Pr(\theta=1)s^1(\Delta_2^2-\Delta_1^2)+\Pr(\theta=0)(1-s^1)(\Delta_2(2-\Delta_2)-\Delta_1(2-\Delta_1))\right)}+1}$$

$$\stackrel{(**)}{\ge} \frac{\frac{(1-\delta)\Delta_1}{\delta+(1-2\delta)\Delta_1}}{\frac{\Delta_1(2\delta+(1-2\delta)\Delta_1)}{\varepsilon(\Delta_2-\Delta_1)((1-2\delta)(\Delta_1+\Delta_2)+2\delta)}+1} \stackrel{(***)}{>} \frac{1+d}{2}.$$

Equality $(*)$ follows from the distribution of signals (1). Observe that the right-hand side of $(*)$ is increasing in s^1. Thus, assuming s^1 is such that (3) holds, we get $(**)$. Finally, by assuming

$$\Delta_1 = \sqrt{\delta}$$

and taking δ close enough to 0, we justify $(***)$:

$$\lim_{\delta\to 0} \left.\frac{\frac{(1-\delta)\Delta_1}{\delta+(1-2\delta)\Delta_1}}{\frac{\Delta_1(2\delta+(1-2\delta)\Delta_1)}{\varepsilon(\Delta_2-\Delta_1)((1-2\delta)(\Delta_1+\Delta_2)+2\delta)}+1}\right|_{\Delta_1=\sqrt{\delta}} = 1.$$

Summing up, we proved

$$\Pr(\theta = i \mid h_t,\, s^i) > \frac{1+d}{2}, \quad i = 0, 1,$$

for δ-extremists. Therefore, the distance between the agents' posteriors exceeds d:

$$\Pr(\theta = 1 \mid h_t,\, s^1) - \Pr(\theta = 1 \mid h_t,\, s^0) = \Pr(\theta = 1 \mid h_t,\, s^1) + \Pr(\theta = 0 \mid h_t,\, s^0) - 1 > 2\frac{1+d}{2} - 1 = d.$$

Proof of Claim 1. We use induction on $k \ge 1$. When agent 0 (1) of level $k = 1$ receives signal $s^0 < \Pr(\theta = 0)$ ($s^1 > \Pr(\theta = 0)$), she considers state i more likely in all periods:

$$\Pr(\theta = 1 \mid s) = \frac{1}{2} \quad \Leftrightarrow \quad \frac{s\Pr(\theta = 1)}{s\Pr(\theta = 1) + (1-s)\Pr(\theta = 0)} = \frac{1}{2} \quad \Leftrightarrow \quad s = \Pr(\theta = 0)$$

Thus, we set x_1^0 to be smaller than $\Pr(\theta = 0)$ and x_1^1 to be larger than $\Pr(\theta = 0)$. For $k > 1$, suppose the claim is true for $k-1$. Without loss of generality, we focus on agent

1 with level k reasoning and signal s^1 that we will restrict later. Consider period $t+1$, with history h_t.

$$
\begin{aligned}
&\Pr(\theta=1 \mid h_t,\, s^1,\, L_0=k-1)\\
&\overset{(1^{**})}{=} \int_0^1 \Pr(\theta=1 \mid h_t,\, s^1,\, s^0,\, L_0=k-1)\Pr(s^0 \mid h_t,\, s^1,\, L_0=k-1)ds^0\\
&\overset{(2^{**})}{=} \int_0^1 \Pr(\theta=1 \mid s^1,\, s^0)\Pr(s^0 \mid h_t,\, s^1,\, L_0=k-1)ds^0\\
&\overset{(3^{**})}{>} \int_0^{x^0_{k-1}} \Pr(\theta=1 \mid s^1,\, s^0)\Pr\left(s^0 \mid s^0 \le x^0_{k-1},\, h_t,\, s^1,\, L_0=k-1\right)ds^0\\
&\overset{(4^{**})}{=} \int_0^{x^0_{k-1}} \Pr(\theta=1 \mid s^1,\, s^0)\frac{\Pr\left(h_t \mid {}^{s^0,\, s^1,}_{L_0=k-1}\right)\Pr\left(s^0 \mid {}^{s^0\le x^0_{k-1},\, s^1,}_{L_0=k-1}\right)}{\Pr\left(h_t \mid {}^{s^0\le x^0_{k-1},\, s^1,}_{L_0=k-1}\right)}ds^0\\
&\overset{(5^{**})}{=} \int_0^{x^0_{k-1}} \Pr(\theta=1 \mid s^1,\, s^0)\Pr\left(s^0 \mid s^1,\, s^0 \le x^0_{k-1}\right)ds^0\\
&\overset{(6^{**})}{=} \Pr\left(\theta=1 \mid s^1,\, s^0 \le x^0_{k-1}\right) \overset{(7^{**})}{\ge} \frac{1}{2}.
\end{aligned}
$$

Equality (1^{**}) holds by the law of total probability. Given signal realizations, the behavior of the agents as well as their levels of reasoning are uncorrelated with the state of the world, which gives us equality (2^{**}).

Fact 1 gives inequality (3^{**}) because conditional probabilities $\Pr(\theta=1 \mid s^1,\, s^0)$ are increasing in s^0.

Equality (4^{**}) follows from Bayes rule. Equality (5^{**}) follows from two observations. First, the signals s^0 and s^1 are independent of the levels of reasoning. Second, by induction hypothesis in all periods $1,\dots,t$ agent 0 of level $k-1$ with signal $s^0 \le x^0_{k-1}$ considers state 0 more likely.

Equality (6^{**}) holds by the law of total probability. By choosing s^1 sufficiently close to 1, we justify (7^{**}). Indeed,

$$
\Pr\left(\theta=1 \mid s^1,\, s^0 < x^0_{k-1}\right) = \frac{s^1 x^0_{k-1}\Pr(\theta=1)}{s^1 x^0_{k-1}\Pr(\theta=1)+(1-s^1)(2-x^0_{k-1})\Pr(\theta=0)},
$$

which approaches 1 as $s^1 \to 1$. That choice gives us x_k^1 and concludes the proof for Claim 1. $\square$

$\square$

As we noted at the end of Section 2, our proof is invariant to how we define an alternative to level k reasoner. The only time we use any assumptions on our agents' beliefs about the opponent's reasoning is inequality (6^*), or more precisely, equality ($*$), and equality ($*$) relies only on the likelihood of level k opponent.

Comment on level updating

Theorem 1 shows that any combination of (d, ε) is possible (the gray square on Figure 1). The key to the result is the growing doubt. Without updating beliefs about the level of reasoning of the opponent, the difference in posteriors of agents who perpetually disagree must be bounded by ε (the hatched triangle on Figure 1). Using $\Pr(L_0 = +\infty)$ instead of $\Pr(L_0 = +\infty \mid h_t,\ s^1)$ and $\Pr(L_0 = k)$ instead of $\Pr(L_0 = k \mid h_t,\ s^1)$, we get

$$\begin{aligned}\Pr(\theta = 1 \mid h_t,\ s^1) = \Pr(L_0 = +\infty)\Pr(\theta = 1 \mid h_t,\ s^1,\ L_0 = +\infty)&\\ + \Pr(L_0 = k)\Pr(\theta = 1 \mid h_t,\ s^1,\ L_0 = k)&\\ \leq (1-\varepsilon)\Pr(\theta = 1 \mid h_t,\ s^1,\ L_0 = +\infty) + \varepsilon.&\end{aligned}$$

Similarly, for agent 0

$$\Pr(\theta = 1 \mid h_t,\ s^0) \geq (1-\varepsilon)\Pr(\theta = 1 \mid h_t,\ s^0,\ L_1 = +\infty).$$

Now, as time passes and the agents continue to disagree, the difference

$$\Pr(\theta = 1 \mid h_t,\ s^1,\ L_0 = +\infty) - \Pr(\theta = 1 \mid h_t,\ s^0,\ L_1 = +\infty)$$

keeps shrinking, and it converges to 0 as t goes to infinite. This is because, under common knowledge of rationality, the agents must "agree in the limit". To see this more concretely, suppose they disagree forever and consider any positive distance between an agent's posterior and $1/2$: after a sufficiently high number of periods of disagreement, the agent realizes that the opponent's disagreement must arise from sufficiently strong evidence in favor of the opposite view, which does not justify the distance of our agent's posterior from $1/2$. Thus, the difference in posteriors is bounded by ε in the limit:

$$\lim_{t\to+\infty} \Pr(\theta = 1 \mid h_t,\ s_t^1) - \Pr(\theta = 1 \mid h_t,\ s_t^0) \leq \varepsilon.$$

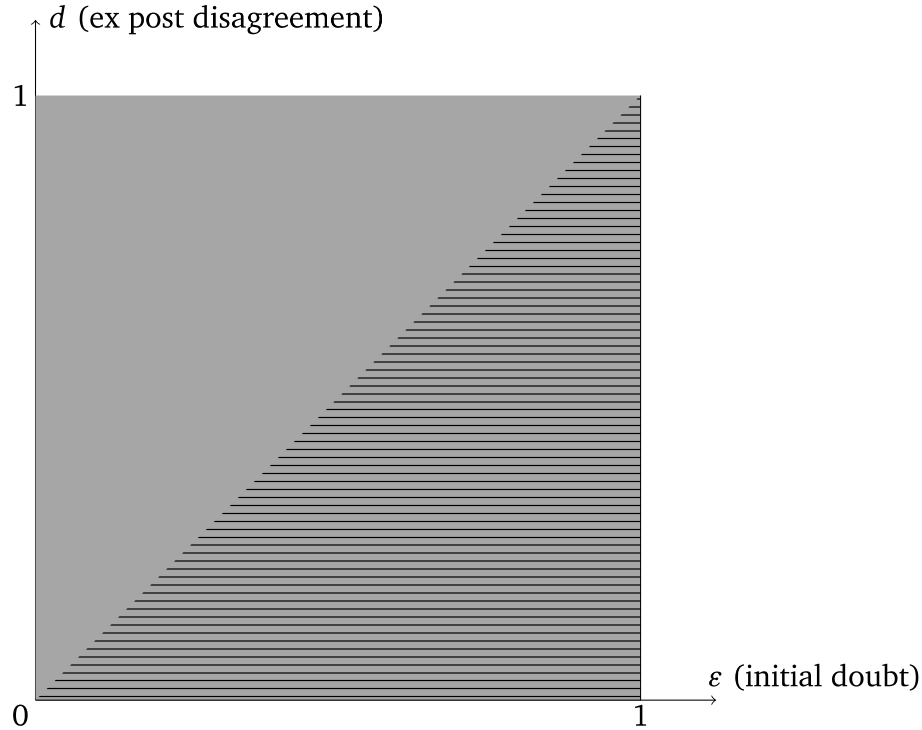

Figure 1: Attainable disagreement with and without level updating. The grey square corresponds to (d,ε) achievable with belief updating that results in growing doubt. The hatched triangle corresponds to (d,ε) achievable when the beliefs about the level of reasoning are not updated.

4 Conclusion

We proposed a model of almost common knowledge of rationality in which large disagreement is possible. It bridges the gap between theoretical impossibility to "agree to disagree" and real life where significant disagreement happens too often to be considered as merely a consequence of some behavioral bias. We show that agents severely disagree ex post if they have some doubt (no matter how small) about an opponent's high level of rationality and their initial disagreement is large enough. Intuitively, higher-order beliefs matter when initial disagreement is large and convergence to a common opinion takes a long time. A doubt placed on higher-order rationality grows over time, disrupting the convergence and sustaining large disagreement.

References

Aumann, Robert J, "Agreeing to disagree," *Annals of Statistics*, 1976, pp. 1236–1239. 2, 3

Bacharach, Michael, "Some extensions of a claim of Aumann in an axiomatic model of knowledge," *Journal of Economic Theory*, 1985, *37* (1), 167 – 190. 3

Bosch-Domenech, Antoni, Jose G Montalvo, Rosemarie Nagel, and Albert Satorra, "One, two,(three), infinity,...: Newspaper and lab beauty-contest experiments," *American Economic Review*, 2002, *92* (5), 1687–1701. 2

Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong, "A cognitive hierarchy model of games," *Quarterly Journal of Economics*, 2004, *119* (3), 861–898. 5

Cave, Jonathan A.K., "Learning to agree," *Economics Letters*, 1983, *12* (2), 147 – 152. 2, 3

DeMarzo, Peter M, Dimitri Vayanos, and Jeffrey Zwiebel, "Persuasion bias, social influence, and unidimensional opinions," *Quarterly Journal of Economics*, 2003, *118* (3), 909–968. 3

Friedenberg, Amanda, Willemien Kets, and Terri Kneeland, "Is bounded rationality driven by limited ability?," 2018. Working Paper. 5

Geanakoplos, John D and Heraklis M Polemarchakis, "We can't disagree forever," *Journal of Economic Theory*, 1982, *28* (1), 192 – 200. 3

Golub, Benjamin and Matthew O Jackson, "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, 2010, *2* (1), 112–149. 3

Hu, Tai-Wei, "On p-rationalizability and approximate common certainty of rationality," *Journal of Economic Theory*, 2007, *136* (1), 379–391. 5

Kneeland, Terri, "Identifying higher-order rationality," *Econometrica*, 2015, *83* (5), 2065–2079. 2

Monderer, Dov and Dov Samet, "Approximating common knowledge with common beliefs," *Games and Economic Behavior*, 1989, *1* (2), 170–190. 3

Morris, Stephen, "Trade with heterogeneous prior beliefs and asymmetric information," *Econometrica*, 1994, pp. 1327–1347. 2

Rubinstein, Ariel, "The Electronic Mail Game: Strategic Behavior Under "Almost Common Knowledge"," *American Economic Review*, 1989, pp. 385–391. 3

Stahl, Dale O. and Paul W. Wilson, "On players' models of other players: Theory and experimental evidence," *Games and Economic Behavior*, jul 1995, *10* (1), 218–254. 4

Strzalecki, Tomasz, "Depth of reasoning and higher order beliefs," *Journal of Economic Behavior & Organization*, 2014, *108*, 108–122. 3

Weinstein, Jonathan and Muhamet Yildiz, "Impact of higher-order uncertainty," *Games and Economic Behavior*, 2007, *60* (1), 200–212. 3