Common Knowledge and State-dependent Equilibria

Nuh Aygun Dalkiran[*]
dalkiran@bilkent.edu.tr

Moshe Hoffman[†]
mohoffman@ucsd.edu

Ramamohan Paturi[‡]
paturi@cs.ucsd.edu

Daniel Ricketts[‡]
daricket@cs.ucsd.edu

Andrea Vattani[‡]
avattani@cs.ucsd.edu

June 2, 2012

Abstract

Many puzzling social behaviors, such as avoiding eye contact, using innuendos, and insignificant events that trigger revolutions, seem to relate to common knowledge and coordination, but the exact relationship has yet to be formalized. Herein, we present such a formalization. We state necessary and sufficient conditions for what we call state-dependent equilibria – equilibria where players play different strategies in different states of the world. In particular, if everybody behaves a certain way (e.g. does not revolt) in the usual state of the world, then in order for players to be able to behave a different way (e.g. revolt) in another state of the world, it is both necessary and sufficient for it to be common p-believed that it is not the usual state of the world, where common p-belief is a relaxation of common knowledge introduced by Monderer and Samet [19]. Our framework applies to many player quorum games – a generalization of coordination games that we introduce – and r-common p-beliefs – a generalization of common p-beliefs that we introduce. We then apply these theorems to three particular signaling structures to obtain novel results. Finally, as a by-product, we resolve some of the outstanding puzzles.

[*]Bilkent University, Department of Economics.
[†]The Rady School of Management & Department of Computer Science and Engineering, University of California San Diego.
[‡]Department of Computer Science and Engineering, University of California San Diego.

1 Introduction

In the popular parable "The Emperor's New Clothes" [2], a gathering of adults pretends to be impressed by the Emperor's dazzling new suit despite the fact that he is actually naked. It is not until an innocent child cries out "But he has nothing on at all!" that the Emperor's position of authority and respect is questioned. This is a metaphor for a number of common political situations in which the populace knows the current regime is inept but takes no action against it until some seemingly insignificant event occurs, such as the child's cry. In fact, in Tunisia, despite years of political repression and poverty, it was not until the previously unknown street vendor Mohamed Bouazizi set himself on fire that citizens rose up in protest. Common knowledge – everyone knows that everyone knows that... – might offer such an explanation for this strange phenomenon: while the boy's cry and the self immolation of Mohamed Bouazizi do not teach anyone that the government is inept, they make it commonly known that the government is inept.

Likewise, common knowledge has been proffered as an explanation for many other puzzling social behaviors: it is common to avoid eye-contact when caught in an inappropriate act, despite the fact that looking away, if anything, increases the conspicuousness of a shameful deed. Nevertheless, even Capuchin monkeys look away when they ignore a request to help an ally in a tussle [25]. And few adults after a nice date are fooled by the inquiry "Would you like to come upstairs for a drink?" yet all but the most audacious avoid the explicit request [26].

Many authors have aptly noted that common knowledge plays an important role in these puzzling social behaviors [11, 4, 7, 26]. Avoiding eye contact prevents common knowledge that you were noticed, using innuendos enables a speaker to request something inappropriate without making the request commonly known, and prohibiting public displays of criticism of the government while not preventing people from realizing the flaws of their government, prevent the flaws from being commonly known. Authors have argued that common knowledge is important in these situations because common knowledge is needed for coordination. But without formal arguments, many important questions still remain, such as: what exactly needs to be "commonly known" in order to "coordinate"? What exactly will happen in the absence of common knowledge? Miscoordination? When common knowledge is lacking, but almost present, e.g. if everyone is pretty sure that everyone is pretty sure... will this have the same effect as common knowledge? Such details, which may seem pedantic, are crucial for answering practical questions such as: if I cannot think up an innuendo, will an appropriately placed cough midsentence do the trick? Why is it that sometimes we use innuendos and sometimes we go out of our way to state the obvious?

We will formalize the role of common knowledge in coordination, which will enable us to address each of these questions in the discussion section. The crucial step in our formalism is based on the insight of Rubinstein [28]. Rubinstein considers coordination games – games in which players make choices such that they would like to mimic the choice that others make. Rubinstein supposes that players coordinate on a particular action A in a given situation. He then supposes that the situation changes and asks whether the players can coordinate on a different action instead. He shows that unless it is commonly known that the situation has changed, players still must coordinate on A. The intuition is clear: even if one player knows that circumstances have changed, if he thinks the other player does not know this, then he expects the other player to play as if circumstances have not changed. Since it is a coordination game, he best responds by playing as if circumstances have not changed. Likewise, even if both players know that circumstances have changed, and both players know that both players know this, but one player does not realize the second player has this second degree of knowledge, then this player will expect the other player to play as if circumstances have not changed. By the above argument, he best responds by playing as if circumstances have not changed. The same logic continues indefinitely.

Rubinstein presents a particular instance in which the above logic holds. The contribution of our paper is to show that this logic holds quite generally, for any two player coordination game, and in fact, for a generalization to many players. And moreover, we show that common knowledge is not just necessary for changing behaviors when circumstances change, but common knowledge is also sufficient. We hope that this will lead to a deeper understanding of these puzzling social behaviors, as well as some novel predictions.

1.1 Our results

In this paper, we introduce *state-dependent equilibria*, which we define as equilibrium strategies in which players take different actions when the circumstances change. This notion allows us to address the questions that were left unanswered by the informal discussions of common knowledge and coordination. In particular, we characterize the conditions under which rational players are able to play state-dependent equilibria.

We begin by considering two-player coordination games. We show that it is not quite common knowledge that determines the existence of state-dependent equilibria but rather a relaxation of common knowledge. This notion corresponds with common p-beliefs, as developed by Monderer and Samet [19]: each believes with probability at least p that each believe with probability at least p.... In our framework, we show that p depends on the precise payoffs of the game and corresponds to the risk dominance of Harsanyi and Selten [15].

We then introduce a natural n-player generalization of coordination games that we call *quorum games* in which coordination on an action is successful if at least some fraction of the players take that action. Accordingly, we also develop a generalization of common p-beliefs for this setting.

In order to derive our results, we provide a unifying theoretical framework for analyzing our games. Our framework gives tight necessary and sufficient conditions on the players' beliefs under which a state-dependent equilibrium exists. These conditions depend on the payoffs of the game (in particular on the risk dominance) and, in the case of quorum games, on the threshold fraction required for successful coordination.

Our final contribution is to apply this framework both to simple but puzzling social behavior and to more complex distributed phenomena that arise in biology, economics, and sociology. The first application is eye-contact. We offer a post hoc explanation for why we avoid eye-contact when caught in an inappropriate act. For the second and third applications, we show how our results can be applied to situations in which the true state of the world is observed by all players with arbitrarily small noise, as in the global games literature [21, 23, 22]. This yields some novel predictions about social behaviors, such as which cues can be used to instigate a revolution, and when a researcher's reputation can be resilient to substandard work. We conclude with a discussion of the social puzzles mentioned in the introduction, suggesting how our theorems can be used to explain these phenomena and what novel prescriptions can be offered.

1.2 Related work

The concept of common knowledge was first formalized in multi-modal logic in 1969 by Lewis [18]. Aumann later put common knowledge in a set-theoretic framework [3].

In 1986, Rubinstein used common knowledge to analyze a problem related to the coordinated attack problem in computer science [28]. This problem, called the Electronic Mail Game, was the first example that common knowledge is very different than any finite order of knowledge. Rubinstein showed that the lack of common knowledge prevents players from switching strategies (i.e. prevents the existence of state-dependent equilibria) in the Electronic Mail Game. See [20] for a retrospective on the Electronic Mail Game. Our results show that common knowledge is not just necessary but also sufficient and holds for any coordination game and even quorum games.

Carlsson and Van Damme showed that when players have noisy signals about the payoffs in a coordination game, as the noise vanishes, the unique equilibrium in the game becomes the risk dominant equilibrium [5]. Morris and Shin applied this result to bank runs and currency crisis, showing that there is a unique underlying value at which currencies collapse and bank runs occur, in contrast to previous models, which permitted multiple equilibria and prevented comparative static analysis [21, 23, 22]. In some of our applications, we use similar signaling structures, but the uncertainty does not affect the payoffs. We find circumstances under which no state-dependent equilibria exist.

Monderer and Samet developed an approximate notion of common knowledge called common p-beliefs, which is relevant in our framework. We will draw heavily on their definitions and results [19].

Others have discussed the role of common knowledge in social puzzles, albeit less formally than in the aforementioned literature. Chwe discusses the role of common knowledge in public rituals [7]. Pinker *et al* discusses the role of common knowledge in innuendos [26]. Binmore and Friedell discuss the role of common knowledge in eye contact [10, 4]. In our paper we formalize the role of common knowledge in many of these social puzzles.

The role of common knowledge has been studied in the fields of distributed computing and artificial intelligence [12, 8, 13]. This line of work suggests that knowledge is an important abstraction for distributed systems and for the design and analysis of distributed protocols, in particular for achieving consistent simultaneous actions. Fagin and Halpern [14, 9] present an abstract model for knowledge and probability in which they assign to each agent-state pair a probability space to be used when computing the probability that a formula is true. A complexity-theoretic version of Aumann's celebrated Agreement Theorem is provided in [1].

2 Preliminaries

We will adopt the set-theoretic formulation of common knowledge introduced by Aumann [3]. In this model, there is a set Ω of "states of the world". Each player i has some information regarding the true state of the world. This information is given by a partition Π_i of Ω. In particular, for $\omega \in \Omega$, $\Pi_i(\omega)$ is the set of states indistinguishable from ω to player i – that is, when ω occurs, player i knows that one of the states in $\Pi_i(\omega)$ occurred but not which one. Finally, there is a probability distribution μ over Ω, representing the (common) prior belief of the players over the states of the world. These parameters all together constitutes the information structure.

Definition 1 (Information structure)**.** *An information structure is a tuple* $\mathcal{I} = (N, \Omega, \mu, \{\Pi_i\}_{i \in N})$ *where*

- N *is the set of players. We let* $n = |N|$.
- Ω *is the set of all possible states of the world.*
- μ *is a strictly positive common prior probability distribution over* Ω.
- Π_i *is the information partition of player* i. $\Pi_i(\omega)$ *gives the set of states indistinguishable from* ω *to player* i.

A (Bayesian) game is now defined by an information structure, a set of possible actions for each player and a state-dependent utility for each player.

Definition 2 (Bayesian game)**.** *A Bayesian game* Γ *is a tuple* $(\mathcal{I}, \{A_i\}_{i \in N}, \{u_i\}_{i \in N})$ *where*

- $\mathcal{I} = (N, \Omega, \mu, \{\Pi_i\}_{i \in N})$ *is an information structure*
- A_i *is the (finite) set of possible actions that player* i *can take.*
- $u_i : A_1 \times A_2 \times \ldots \times A_n \times \Omega \to \mathcal{R}$ *is the utility for player* i *given the state of the world and the actions of all players.*

A strategy profile prescribes the action (possibly randomized) that each player takes at each state of the world.

Definition 3 (Strategy profile)**.** *A strategy profile is a function* $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_n) : \Omega \to A_1 \times A_2 \times \ldots \times A_n$ *that specifies what action each player take in each state of the world.*

Since a player cannot distinguish between states belonging to the same partition, it is enforced that if a player i plays some strategy $\sigma = \sigma_i(\omega)$ at some state $\omega \in \Omega$, it must be the case that i plays σ at all states $\omega' \in \Pi_i(\omega)$. We can now recall the definition of Bayesian Nash equilibrium.

Definition 4 (Bayesian Nash equilibrium)**.** *A strategy profile* $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_n) : \Omega \to A_1 \times A_2 \times \ldots \times A_n$ *is a Bayesian Nash equilibrium (BNE) of* Γ *if for all* $i \in N$,

1. $\sigma_i(\omega) = \sigma_i(\omega')$ *whenever* $\omega \in \Pi_i(\omega')$.
2. $\int_{\omega \in \Omega} u_i(\sigma_i(\omega), \sigma_{-i}(\omega)) \mathrm{d}\mu(\omega) \geq \int_{\omega \in \Omega} u_i(\sigma'_i(\omega), \sigma_{-i}(\omega)) \mathrm{d}\mu(\omega)$ *for all* σ' *satisfying property 1.*

2.1 Beliefs

We now define the notion of p-belief, introduced by Monderer and Samet [19], which extends the notion of common knowledge by Aumann [3]. Let p be a number between 0 and 1. We say that a player i p-believes the event E at state of the world ω if the subjective probability that i assigns to E at ω is at least p. That is, whenever ω is the true state of the world, i believes that an event in E occurred with probability at least p. Henceforth, we will use short expressions such as "i p-believes E at ω" to refer to this concept.

We denote by $B_i^p(E)$ the set of all states of the world at which player i p-believes E.

Definition 5 (p-belief)**.** *For any* $0 \leq p \leq 1$, *we say that player* i p*-believes* E *at* ω *if* $\mu(E \mid \Pi_i(\omega)) \geq p$. *We will denote by* $B_i^p(E)$ *the event that* i p*-believes* E, *i.e.*

$$B_i^p(E) = \{\omega \mid \mu(E \mid \Pi_i(\omega)) \geq p\}.$$

Observe that by definition of $B_i^p(E)$, the notation $\omega \in B_i^p(E)$ indicates that whenever ω occurs, player i believes with probability at least p that the event E occurred. An event E is then defined p-evident if whenever it occurs, each player i believes with probability at least p that it indeed occurred.

Definition 6 (evident p-belief)**.** *An event* E *is evident* p*-belief if for all* $i \in N$ *we have* $E \subseteq B_i^p(E)$.

The following concept extends the notion of common knowledge.

Definition 7 (common p-belief)**.** *An event* C *is common* p*-belief at state* ω *if there exists an evident* p*-belief event* E *such that* $\omega \in E$, *and for all* $i \in N$, $E \subseteq B_i^p(C)$.

Note that an evident p-belief event E is trivially common p-believed at any state $\omega \in E$. That is, if a state $\omega \in E$ occurs, the occurrence of E is p-believed by all players.

An alternative (non-fixed point) characterization of a common p-belief event is the following. We refer to [19] for its proof.

Lemma 1 (common p-belief characterization)**.** *Given an event C, let $C^0 = C$ and inductively define $C^n = \bigcap_i B_i^p(C^{n-1})$. Then C is common p-belief at state ω if and only if $\omega \in \bigcap_{n\geq 1} C^n$.*

2.2 Game setup

We will focus on a generalization of coordination games in which there are n players, each of whom can take action A or B. A player's payoffs are a function of whether the fraction of players who play a specific strategy exceeds a threshold $\bar{r}$. When $n = 2$ and the threshold $\bar{r} = 1$, we obtain the classic 2-player coordination game. The precise definition for the case $n > 2$ is provided in Section 5.

Definition 8 (State-dependent BNE)**.** *We say that a Bayesian Nash equilibrium σ^* is state-dependent if for some $\omega, \omega' \in \Omega$, $i \in N$, we have that $\sigma_i^*(\omega) = A$ and $\sigma_i^*(\omega') = B$.*

Note that in a 2-player game an equilbrium σ^* is state-dependent if and only if there exist $\omega, \omega' \in \Omega$ such that $\sigma^*(\omega) = (A, A)$ and $\sigma^*(\omega') = (B, B)$.

3 Two player framework

In this section we consider the classic 2-player, 2-strategy symmetric coordination game. The payoffs are as follows:

	A	B
A	a, a	b, c
B	c, b	d, d

Assumption 1 (Coordination game)**.** *We make the following standard assumption on the parameters of a symmetric coordination game: $a > c$ and $d > b$.*

Throughout this paper, we will use $p^* = \frac{d-b}{d-b+a-c}$. This value is called risk-dominance [15]. Note that if player i believes with probability exactly p that the other player will play A at ω, then player i will be indifferent between playing A and B at ω.

For convenience, we will use the following definitions throughout this section.

Definition 9. *Given any strategy profile σ, we let $A_i(\sigma) = \{\omega | \sigma_i(\omega) = A\}$ and $B_i(\sigma) = \{\omega | \sigma_i(\omega) = B\}$, i.e. the set of states where player i plays A and B respectively.*

We now state our main result for the 2-player case. The main question we ask is when it is possible for the two players to coordinate on different actions in different states of the world. We answer this question in terms of the existence of evident p-belief events (where p depends on the payoff matrix) showing that such events are necessary and sufficient.

Theorem 1. *There exists a state-dependent Bayesian Nash equilibrium σ^* if and only if there exists a non-empty evident p^*-belief event E and a non-empty evident $(1-p^*)$-belief event F such that $E \cap F = \emptyset$.*

Proof. First we will show that if there exists a state-dependent Bayesian Nash equilibrium σ^*, then there exists a non-empty evident p^*-belief event E and a non-empty evident $(1-p^*)$-belief event F such that $E \cap F = \emptyset$. In particular, we will show that $A_1(\sigma^*) \cap A_2(\sigma^*)$ is evident p^*-belief and $B_1(\sigma^*) \cap B_2(\sigma^*)$ is evident $(1-p^*)$-belief. In the rest of the proof, we will let $A_i = A_i(\sigma^*)$ and $B_i = B_i(\sigma^*)$.

We need to show that $A_1 \cap A_2$ and $B_1 \cap B_2$ are both non-empty. This follows from the fact that σ^* is state-dependent; if a player plays A on some state, then he must believe that there is some state in which both players choose A. If this were not true, then he would best respond by playing B.

We also must show that $A_1 \cap A_2 \cap B_1 \cap B_2 = \emptyset$. This follows from the fact that $A_1 \cap B_1 = \emptyset$.

Assume for contradiction that $A_1 \cap A_2$ is not evident p^*-belief, and without loss of generality, assume that $A_1 \cap A_2 \not\subseteq B_1^{p^*}(A_1 \cap A_2)$. Consider some $\omega \in A_1 \cap A_2$ such that $\omega \notin B_1^{p^*}(A_1 \cap A_2)$. It must be the case that $\mu_1(A_1 \cap A_2 | \Pi_1(\omega)) < p^*$. Furthermore, by definition 4.1, $\Pi_1(\omega) \subseteq A_1$. Therefore, we must have $\mu_1(A_2 | \Pi_1(\omega)) <$

p^*, i.e. at ω, player 1 believes with probability strictly less than p^* that player 2 will play A. Thus, with probability strictly less than p^*, player 1 will get payoff a if he plays A and payoff c if he plays B. Furthermore, with probability at least $1-p^*$, he will get payoff b if he plays A and payoff d if he plays B. Since $a-c \geq b-d$ (from assumption 1), we obtain

$$u_i(A|\Pi_i(\omega)) - u_i(B|\Pi_i(\omega)) < p^*(a-c) + (1-p^*)(b-d) = 0 \quad (1)$$

Equation (1) implies that σ^* is not a Bayesian-Nash equilibrium, and hence we obtain a contradiction. The proof that F is evident $(1-p^*)$-belief is analogous.

Now we will prove that if there exists a non-empty evident p^*-belief event E and a non-empty evident $(1-p^*)$-belief event F such that $E \cap F = \emptyset$, then there exists a state-dependent Bayesian-Nash equilibrium σ^*. We will do so by constructing such a Bayesian-Nash equilibrium.

First, we will recursively define the sequence of events E_i^n and F_i^n as follows. For all i, let $E_i^0 = E$ and $F_i^0 = F$. For all i and for all $n \geq 1$, let $E_i^n = B_i^{p^*}(E_{-i}^{n-1})$ and $F_i^n = \bigcup_{q<p^*} B_i^{1-q}(F_{-i}^{n-1})$. Finally, let $\hat{E}_i = \bigcup_{n=0}^{\infty} E_i^n$ and $\hat{F}_i = \bigcup_{n=0}^{\infty} F_i^n$.

Now let σ^* be a strategy profile such that $\sigma_i^*(\omega) = A$ for all $\omega \in \hat{E}_i$, $\sigma_i^*(\omega) = B$ for all $\omega \in \hat{F}_i$, and $\sigma_i^*(\omega) = B$ for all $\omega \in \hat{E}_i^{\,c} \cap \hat{F}_i^{\,c}$. We will show that σ^* is a Bayesian Nash equilibrium using the following 4 steps:

1. $E_i^n \subseteq E_i^{n+1}$ and $F_i^n \subseteq F_i^{n+1}$.
2. $E_i^n \cap F_i^n = \emptyset$ for all $n \geq 0$.
3. For all $\omega \in \hat{E}_i$, player i will be no better off by playing B.
4. For all $\omega \in \hat{E}_i^{\,c} \cap \hat{F}_i^{\,c}$, player i will be no better off by playing A.

The first two steps show that σ^* is well defined. The final two steps, combined with an analogous argument for all $\omega \in \hat{F}_i$ show that σ^* is a Bayesian Nash equilibrium.

Proof of step 1: We will show that $E_i^n \subseteq E_i^{n+1}$ by induction on n. For the base case, note that $E_i^0 = E \subseteq B_i^{p^*}(E) = E_i^1$ for all i because E is an evident p^*-belief event. Now suppose $E_i^{n-1} \subseteq E_i^n$ for all i. Then by monotonicity of p-beliefs,[1] $E_i^n = B_i^{p^*}(E_{-i}^{n-1}) \subseteq B_i^{p^*}(E_{-i}^n) = E_i^{n+1}$. An analogous argument proves that $F_i^n \subseteq F_i^{n+1}$.

Proof of step 2: We will show that $E_i^n \cap F_i^n = \emptyset$ for all $n \geq 0$ by induction on n. For the base case, note that $E_i^0 \cap F_i^0 = E \cap F = \emptyset$. Now suppose $E_i^n \cap F_i^n = \emptyset$ for all i. Then $\mu(E_{-i}^n \mid \Pi_i(\omega)) + \mu(F_{-i}^n \mid \Pi_i(\omega)) \leq 1$ for all ω and for all i. If $\omega \in E_i^{n+1}$, then by definition of E_i^{n+1}, $\mu(E_{-i}^n \mid \Pi_i(\omega)) \geq p^*$. This implies that $\mu(F_{-i}^n \mid \Pi_i(\omega)) \leq 1-p^*$. Therefore, $\omega \notin F_i^{n+1}$.

Proof of step 3: Suppose $\omega \in \hat{E}_i$. Then by the construction of $\hat{E}_i$, $\omega \in E_i^n$ for some $n \geq 1$ or $\omega \in E$. First consider the case in which $\omega \in E$. Since E is evident p^*-belief, $\mu(E \mid \Pi_i(\omega)) \geq p^*$. Since player $-i$ plays A on E, player i can do no better by playing B at ω.

Now consider the case in which $\omega \in E_i^n$ for some $n \geq 1$. In this case, $\mu(E_{-i}^{n-1} \mid \Pi_i(\omega)) \geq p^*$. Since player $-i$ plays A on E_{-i}^{n-1}, player i can do no better by playing B at ω.

Proof of step 4: Suppose $\omega \in \hat{E}_i^{\,c} \cap \hat{F}_i^{\,c}$. Then for all n, $\omega \notin E_i^n$. Therefore, there is no n such that $\mu(E_{-i}^{n-1} \mid \Pi_i(\omega)) \geq p^*$. Since player $-i$ only plays A on $\hat{E}_i$, this means that at ω, player i believes with probability strictly less than p^* that player $-i$ will play A. Therefore, player i will be no better off by playing A at ω. □

While evident knowledge is both necessary and sufficient for state-dependent equilibria, our theorem further allows us to specify how the strategies must depend on these evident events, which we express in the following corollary:

Corollary 1. *A strategy profile σ^* is a state-dependent Bayesian Nash equilibrium if and only if there exists a non-empty evident p^*-belief event E and a non-empty evident $(1-p^*)$-belief event F such that $B_i^{p^*}(E) \cap B_i^{1-p^*}(F) = \emptyset$ and $B_i^{p^*}(E) \cup B_i^{1-p^*}(F) = \Omega$ for all i, in which case $A_i(\sigma^*) = B_i^{p^*}(E)$ and $B_i(\sigma^*) = B_i^{1-p^*}(F)$ for all i.*

Our next corollary states the relationship between state-dependent equilibria and common knowledge.

Corollary 2. *If σ^* is a Bayesian Nash equilibrium such that $\sigma_i^*(\omega) = A$ and $\sigma_i^*(\omega') = B$, then $\neg\omega'$ is common p^*-belief at ω and $\neg\omega$ is common $(1-p^*)$-belief at ω'.*

[1] If $F \subseteq G$, then $B_i^p(F) \subseteq B_i^p(G)$ for any i and any p [19].

4 Two player applications

4.1 A Rationale for Avoiding Eye-Contact

In this section, we use corollary 2 to provide a rationale for avoiding eye-contact when one has committed a socially deviant act. The basic intuition is that avoiding eye-contact prevents switching from a more desirable equilibrium to a less desirable equilibrium.

THE STORY. Two Charedi men, Bob and Dave, go to a bar, and each spots the other, purposely looking away before meeting eyes. Why?

Suppose that the next day they have to decide whether to tell the Rabbi. If one expects the other to tell, he is better off also admitting to his actions. On the other hand, if one does not expect the other to tell, then he is better off also not admitting to his transgression. The payoffs can be interpreted as the coordination game from the two-player framework by interpreting A as the act of not telling the Rabbi, B as the act of telling the Rabbi.

We make the reasonable assumption that if at least one of the men stays home, neither tells the Rabbi that he saw the other player at the bar (since he in fact did not).

We will use our framework from section 3 to show that (a) there is always an equilibrium in which they both tell the Rabbi if they make eye-contact at the bar, and (b) under mild assumptions, if they do not make eye contact, neither will tell the Rabbi.

THE MODEL. We now specify the information structure: we suppose that in one state of the world, at least one of them stays home ($\mathcal{H}$) while in another state of the world, Dave enters the bar, and Bob is already sitting at the bar. When Dave walks in, Bob is either staring at the bartender, in which case he would not see Dave, or looking at the door, in which case he would. As soon as Dave enters, he sees Bob, so he quickly turns around and walks out. Dave turns around before or after noticing if Bob saw him.

The set of possible states of the world is given by

$$\Omega = \{\mathcal{H}, (\mathcal{R}, \mathcal{G}), (\mathcal{R}, \mathcal{G}'), (\mathcal{R}', \mathcal{G}), (\mathcal{R}', \mathcal{G}')\}.$$

We interpret the states of the world as follows: $\mathcal{H}$ is the state where Dave *does not* go to the bar and stays at ($\mathcal{H}$)ome. $(\mathcal{R}, \mathcal{G})$ is the state where Dave goes to the bar, Bob is looking at the bartender, and Dave leaves the bar before checking if Bob saw him. $(\mathcal{R}', \mathcal{G}')$ is the state where Dave goes to the bar, Bob sees him, and Dave sees that Bob saw him (i.e. they make eye-contact).

The information partitions are given as follows:

- $\Pi_A = \{\{\mathcal{H}, (\mathcal{R}, \mathcal{G}), (\mathcal{R}, \mathcal{G}')\}, \{(\mathcal{R}', \mathcal{G})\}, \{(\mathcal{R}', \mathcal{G}')\}\}$
- $\Pi_M = \{\{\mathcal{H}\}, \{(\mathcal{R}, \mathcal{G}), (\mathcal{R}', \mathcal{G})\}, \{(\mathcal{R}, \mathcal{G}')\}, \{(\mathcal{R}', \mathcal{G}')\}\}$

Observe that $(\mathcal{R}', \mathcal{G}')$ is an evident p^*-belief event, that is, when eye contact happens, it becomes common knowledge between Bob and Dave as expected.

We use the following independent probabilities to deduce the priors over the state space: p_b is the probability that Dave goes to the bar i.e., he *does not* stay home; p_r is the probability that Bob is looking at the bartender when Dave walks in; p_g is the probability that, conditioned on Dave going to the bar, he leaves the bar before checking if Bob saw him.

Our first claim is an almost trivial one which shows that there always exists an equilibrium in which they both tell the Rabbi if they make eye-contact at the bar.

Claim 1. *There exists a Bayesian-Nash equilibrium of* Γ *such that* $\sigma^*(\mathcal{H}) = (A, A)$ *and* $\sigma^*((\mathcal{R}', \mathcal{G}')) = (B, B)$ *for any* p_b, p_r, p_g.

Proof. Let $\sigma^*(\omega) = (A, A)$ for all $\omega \notin \Pi_i((\mathcal{R}', \mathcal{G}'))$ for both i. The sufficiency conditions for Corollary 1 is then met by letting $F = \{(\mathcal{R}', \mathcal{G}')\}$ and $E = \Omega \backslash F$. □

Our next claim shows conditions under which if Bob and Dave do not make eye-contact, they must continue playing A if they play A on $\mathcal{H}$. That is, suppose Bob and Dave coordinate on (A, A) when Dave stays home; under what conditions is it the case that they can play (B, B) only at $(\mathcal{R}', \mathcal{G}')$, i.e. only when they make eye-contact. We defer the proof to appendix A.1.

Claim 2. *Suppose σ^* is a Bayesian-Nash equilibrium of Γ with $\sigma^*(\mathcal{H}) = (A, A)$. If $p_r > p^*$ and $\frac{p_b p_r}{p_b p_r + (1-p_b)} < 1-p^*$ then $\sigma^*(\omega) \neq (B, B)$ for all $\omega \in \Omega \backslash \{(\mathcal{R}', \mathcal{G}')\}$.*

Now that we have formalized why someone might want to avoid eye contact, we can discuss when this is worthwhile. For instance, avoiding eye contact will not serve any purpose when it is very likely that they saw each other, e.g. if the bar had nobody else present and was very well lit (i.e. when p_r and p_g are small). Likewise, avoiding eye contact serves no purpose if, when it is commonly known that both parties see each other doing an act, neither is expected to play any differently than if neither transgressed (i.e. $\sigma^*(\mathcal{R}', \mathcal{G}') = (A, A)$). For example, the transgression is not perceived as related to the ensuing coordination game, e.g. if the two religious men have already discussed their secret abhorrence of the religion.

Moreover, avoiding eye contact only serves a purpose if there will be an ensuing coordination game (i.e. $a > c$). If in fact Bob would prefer to rat on Dave, regardless of whether Dave rats on Bob (e.g. because he knows the Rabbi will believe him, and he would like Dave to be excommunicated) then Dave does not help himself by avoiding Bob's eyes. In fact, to the extent that Dave thinks this might be the case, he might want to avoid eye contact, as it may make his presence more conspicuous to Bob.

Lastly, Bob may even purposely make eye contact, or yell out "hey Dave, is that you," if he in fact *wants* to switch from them both playing A to both playing B (which would be the case if $d > a$). For instance, this would be the case if Dave was looking for someone to leave the community with him and help him start a new life in the secular world.

5 n-player framework

5.1 Game setup

Let Ω be all possible states of the world. There are n players, each of whom can take action A or B. A player's payoff for a particular action is a function of the fraction of players who play B. In particular, a player's payoffs are a function of whether the fraction of players who play B exceeds a threshold $\bar{r}$. Let r denote the fraction of players who play B. The payoffs are as follows.

$$u_i(A, r) = \begin{cases} a & : r \leq \bar{r} \\ b & : r > \bar{r} \end{cases} \qquad u_i(B, r) = \begin{cases} c & : r \leq \bar{r} \\ d & : r > \bar{r} \end{cases}$$

We again use assumption 1 on the values of the parameters, namely that $a > c$ and $d > b$. In this context, these assumptions on the payoff parameters generalize that of a 2 player coordination game in that a player best respond by playing A if and only if sufficiently many others play A.

We will also assume that n is sufficiently large such that a particular player's decision to play A or B does not affect whether r exceeds $\bar{r}$.

Furthermore, we will again use $p^* = \frac{d-b}{d-b+a-c}$. For n-players, p^* is a generalization of risk dominance. If player i believes with probability exactly p that at least $(1 - \bar{r})$ players will play A at ω, then player i will be indifferent between playing A and B at ω.

Note that this setup is a generalization of the two player setup. In particular, if there are two players, then we can let $\bar{r}$ be any value in $(0, 1)$ in order to obtain the two player model.

5.2 Main Result

Definitions 10, 11, and 12 generalize p-beliefs and evident p-beliefs, and common p-beliefs to n players.

Definition 10 (r, p-belief)**.** *For any $0 \leq p \leq 1$ and any $0 \leq r \leq 1$, we say that event E is r, p-belief at ω if $|\{i \mid \omega \in B_i^p(E)\}| \geq rn$. We will denote by $B^{r,p}(E)$ the event that at least a fraction of r players p-believes E, i.e.*

$$B^{r,p}(E) = \{\omega \mid |\{i \mid \omega \in B_i^p(E)\}| \geq rn\}$$

The following lemma extends a property of p-beliefs, proved in [19], to r, p-beliefs.

Lemma 2 (Monotonicity of r, p-beliefs)**.** *If $F \subseteq G$, then $B^{r,p}(F) \subseteq B^{r,p}(F)$.*

Proof. This follows from monotonicity of p-beliefs [19]. ∎

Definition 11 (r-evident p-belief)**.** *An event E is r-evident p-belief if $E \subseteq B^{r,p}(E)$.*

Definition 12 (r-common p-belief). *Given an event C, let $C^0 = B^{r,p}(C)$ and inductively define $C^n = B^{r,p}(\bigcap_{i<n} C^i)$ for all $n \geq 2$. Then C is r-common p-belief at ω if $\omega \in \bigcap_{n\geq 1} C^n$*

While r-common p-beliefs is identical to common p-beliefs when $n = 2$ and $r = 1$, it should be noted that there are multiple ways one could have extended common p-beliefs to n players that are not all equivalent. Definition 12 gives the only way of extending common p-beliefs such that they relate to r-evident p-beliefs, as described in the following lemma. The proof is deferred to appendix A.2.

Lemma 3 (r-common p-belief fixed point characterization). *C is r-common p-belief at ω if and only if there exists an r-evident p-belief event E such that $\omega \in E$ and $E \subseteq B^{r,p}(C)$.*

The following theorem and corollaries are analogous to our two-player theorems and corollaries, despite the differing setup and proofs. The proof of theorem 2 is provided in appendix A.3.

Theorem 2. *There exists a state-dependent Bayesian Nash equilibrium σ^* if and only if there exists a non-empty $(1-\bar{r})$-evident p^*-belief event E and a non-empty s-evident $(1-p^*)$-belief event F such that $E \cap F = \emptyset$ for some $s > \bar{r}$.*

Corollary 3. *A strategy profile σ^* is a state-dependent Bayesian Nash equilibrium if and only if there exists a non-empty $(1-\bar{r})$-evident p^*-belief event E and a non-empty s-evident $(1-p^*)$-belief event F for some $s > \bar{r}$ such that $B_i^{p^*}(E) \cap B_i^{1-p^*}(F) = \emptyset$ and $B_i^{p^*}(E) \cup B_i^{1-p^*}(F) = \Omega$ for all i, in which case $A_i(\sigma^*) = B_i^{p^*}(E)$ and $B_i(\sigma^*) = B_i^{1-p^*}(F)$ for all i.*

Corollary 4. *If σ^* is a Bayesian Nash equilibrium such that $|\{j \mid \sigma_j^*(\omega) = A\}| \geq 1-\bar{r}$ and $|\{j \mid \sigma_j^*(\omega') = B\}| > \bar{r}$, then $\neg\omega'$ is $(1-\bar{r})$-common p^*-belief at ω and $\neg\omega$ is $\bar{r}$-common $(1-p^*)$-belief at ω'.*

6 n-player applications

6.1 Tea Taxes

THE STORY. Before the American Revolution, the British government raised and lowered existing taxes on numerous occasions without instigating a revolution, even when the taxes were high. However, as soon as the British introduced a new tax on tea, then the colonists revolted, even though the tax was lower than before [6]. Why?

Clearly if too few people revolt, then the British empire can easily quash the rebellion and punish those who participated. However, if sufficiently many revolt, the British cannot quash the rebellion and those who participated get the rewards of a new government, with the possibility of punishing the "benedict arnolds" who did not participate. This can be represented as a quorum game by interpreting A as the act of not revolting, B as the act of revolting, and $\bar{r}$ as the number of revolters needed to overwhelm the British.

THE MODEL. We model the gradual adjustments of the existing tax as follows: the true tax rate is ω, which is uniformly distributed between 0 and 1. We assume that the players do not perfectly observe the tax rate; instead each colonist i observes a noisy signal s_i of the true tax rate drawn independently from $\mathcal{U}[\omega-\epsilon, \omega+\epsilon]$. Possible reasons why they might not exactly know the tax rate include that each individual does not know how much of the various taxed goods they will consume, they do not know how perfectly the taxes will be enforced, and it is hard to keep up with the exact taxes since they change fairly often. Our results are robust to any amount of noise $\epsilon > 0$. This signal structure is reminiscent of the global games literature [21, 23, 22].

We contrast the gradual adjustments of the existing tax with the discrete signal from the tea tax. We model this scenario as follows. The government creates a new tax only if their total tax rate is above a certain threshold, i.e. if $\omega > M$ for some M. However, the colonists only observe whether there is a new tax with some noise. To make this situation comparable to the previous, we will assume that the colonists have strictly less information. In particular, we will assume that the discrete signal is produced based on the same signaling structure as above, but instead of observing s_i, player i observes only h_i, which is 0 if s_i is below M and 1 otherwise. If a new tax is rumored to be introduced but is not actually introduced, there is a small probability that some individual might wrongly think the new tax was introduced (i.e. $\omega < M$, but $s_i > M$, hence $h_i = 1$).

Using our n-player framework, we will show that it is not possible for the colonists to revolt if the existing taxes increase gradually, but it is possible for them to revolt when a new tax is introduced.[2]

As in the general setup, the payoff from each action is a function of the fraction of players who play B. Let r denote the fraction of players who play B. The payoffs are as follows.

$$u_i(A,r) = \left\{ \begin{array}{ll} a & : r \leq \bar{r} \\ b & : r > \bar{r} \end{array} \right. \qquad u_i(B,r) = \left\{ \begin{array}{ll} c & : r \leq \bar{r} \\ d & : r > \bar{r} \end{array} \right.$$

6.1.1 Gradual tax increase

We first consider the scenario in which each player observes a private continuous noisy signal. In this scenario, player i observes a private signal s_i that is drawn from $\mathcal{U}[\omega - \epsilon, \omega + \epsilon]$. We want to show conditions under which it is not possible for players to coordinate on whether their signal is above or below a particular value. Theorem 3 gives such a result.

Theorem 3 (Continuous signals)**.** *Let $\omega \sim \mathcal{U}[0,1]$ and $s_i \sim^{iid} \mathcal{U}[\omega - \epsilon, \omega + \epsilon]$ for all i and for $\epsilon > 0$. Let σ^* be a strategy profile such that $\sigma_i^*(s_i) = B$ when $s_i \geq s^*$ and $\sigma_i^*(s_i) = A$ when $s_i < s^*$ for some $s^* \in [\epsilon, 1-\epsilon]$. Then σ^* is not a Bayesian Nash equilibrium except in the knife edge case in which $\bar{r} = p^*$.*

Proof. We will use corollary 3 to demonstrate that no Bayesian Nash equilibrium exists of the form specified in this theorem. Let E be the set of states in which at least $(1-\bar{r})$ fraction of the players choose A and let $F = E^c$, i.e. the set of states in which at least $\bar{r}$ fraction of the players choose B. By corollary 3, every state in E must be p^*-believed by at least $(1-\bar{r})$ fraction of the players, and every state in F must be $(1-p^*)$-believed by at least $\bar{r}$ fraction of the players. We will show that this is true only if $\bar{r} = p^*$.

First, observe that if the true state is ω and $\omega \in [s^* - \epsilon, s^* + \epsilon]$, then the fraction of players who receive a signal above s^* (and thus play B) is given by $\frac{\omega + \epsilon - s^*}{2\epsilon}$. For $\omega \geq s^* + \epsilon$, the fraction of players who play B is 1 and for $\omega \leq s^* - \epsilon$, the fraction of players who play B is 0. Therefore, there exists some threshold ω^*, above which at least $\bar{r}$ fraction of the players choose B and below which at least $(1-\bar{r})$ fraction of the players choose A. In particular, $\omega^* = 2\epsilon\bar{r} - \epsilon + s^*$, $F = \{\omega \mid \omega \geq \omega^*\}$, and $E = \{\omega \mid \omega < \omega^*\}$.

It must be the case that for any $\omega \in F$, at least $\bar{r}$ fraction of the players will receive private signals that induce a posterior probability of at least $(1-p^*)$ that $\omega \geq \omega^*$. For any player i with private signal $s_i \in [\omega^* - \epsilon, \omega^* + \epsilon]$, the posterior probability that $\omega \geq \omega^*$ is $\frac{s + \epsilon - \omega^*}{2\epsilon}$. Thus, in order for player i to have posterior probability at least $(1-p^*)$ that $\omega \geq \omega^*$, his private signal must satisfy the condition $\frac{s+\epsilon-\omega^*}{2\epsilon} \geq 1 - p^*$. Thus, it must be the case that

$$s_i \geq \epsilon - 2\epsilon p^* + \omega^* \tag{2}$$

The fraction of players who receive such a signal is increasing in ω. This implies that if at least $\bar{r}$ fraction of the players receive a signal satisfying equation (2) at ω^*, then at least $\bar{r}$ fraction receive such a signal for any $\omega \geq \omega^*$. When $\omega = \omega^*$, the fraction of players who receive a private signal satisfying equation (2) is given by

$$\frac{\omega^* + \epsilon - (\epsilon - 2\epsilon p + \omega^*)}{2\epsilon} = p^* \tag{3}$$

This means that at least $\bar{r}$ fraction of the players will $(1-p^*)$-believe F at every state in F if and only if $p^* \geq \bar{r}$.

However, it must also be the case that for any $\omega \in E$, at least $(1-\bar{r})$ fraction of the players will receive private signals that induce a posterior probability of at least p^* that $\omega < \omega^*$. By an analogous argument, the fraction of players who receive such a signal as ω approaches ω^* is $(1p)$. Therefore, σ^* is a Bayesian Nash equilibrium only if $1 - p^* \geq 1 - \bar{r}$ and $p^* \geq \bar{r}$, which happens only when $p^* = \bar{r}$. □

6.1.2 Introduction of new tax

We now contrast the previous result with a scenario in which signals are discrete, rather than continuous. As mentioned before, we apply a filter to each private signal s_i before player i observes it. Each s_i will be drawn independently from a distribution that is a function of ω. However, player i will not observe s_i. Instead, player i will observe a $h_i = 0$ if $s_i < M$ and $h_i = 1$ if $s_i \geq M$. In other words, each player observes whether his signal is less than or greater than some commonly known threshold M, but does not observe the actual value of his signal.

[2]We imagine that the taxes do not effect the payoffs in the quorum game. In this particular example, this is somewhat reasonable given that the tax rate under consideration was much smaller than the amount spent on the war effort or the gains or losses from winning or losing the war. Nonetheless, our results are robust to state dependent payoffs, so long as the state does not have too dire an impact on the payoffs.

Theorem 4 (Binary signals)**.** *Let* $\omega \sim \mathcal{U}[0,1]$ *and* $s_i \sim^{iid} \mathcal{U}[\omega-\epsilon, \omega+\epsilon]$ *for all* i*. Then there exists a Bayesian Nash equilibrium* σ^* *in which* $\sigma_i^*(0) = A$ *and* $\sigma_i^*(1) = B$ *when* $M \in [\epsilon, 1-\epsilon]$*, and* $p^* \in [\frac{\epsilon}{1-M}, 1-\frac{\epsilon}{M}]$*.*

Proof. We will use corollary 3 to construct the appropriate Bayesian Nash equilibrium for this theorem. We will let the state space Ω be the set of binary vectors of length n, where the ith element is the private signal of player i. Player i can only distinguish vectors in which the ith element is a 0 from vectors in which the ith element is a 1.

By corollary 3, it is sufficient to demonstrate two events E and F such that at least $(1-\bar{r})$ fraction of the players p^*-believe E at every state in E, strictly more than $\bar{r}$ fraction of the players $(1-p^*)$-believe F at every state in F, $B_i^{p^*}(E) \cap B_i^{1-p^*}(F) = \emptyset$, and $B_i^{p^*}(E) \cup B_i^{1-p^*}(F) = \Omega$. We let $E = \{0^n\}$ and $F = \{1^n\}$.

First we will show that $F \subseteq B^{\bar{r},1-p^*}(F)$. Since $h_i = 1$ at F for all i, all players $(1-p^*)$-believe F at F if $\Pr(F \mid h_i = 1) = \Pr(\omega \geq M+\epsilon \mid h_i = 1) \geq 1-p^*$. Since the true state and noise are both uniform,

$$\Pr(\omega \geq M+\epsilon \mid h_i = 1) = \frac{1-(M+\epsilon)}{1-(M+\epsilon)+\epsilon} \geq 1-p^* \tag{4}$$

By an analogous argument, every player p^*-believes E at E since

$$Pr(\omega < M-\epsilon \mid h_i = 0) = \frac{M-\epsilon}{M} \geq p^* \tag{5}$$

Furthermore, $\Pr(E \mid h_i = 1) = 0$ and $\Pr(F \mid h_i = 0) = 0$. Thus, player i p^*-believes E exactly when $h_i = 0$ and $(1-p^*)$-believes F exactly when $h_i = 1$. Therefore, $B_i^{p^*}(E) \cap B_i^{1-p^*}(F) = \emptyset$, and $B_i^{p^*}(E) \cup B_i^{1-p^*}(F) = \Omega$.

Finally, corollary 3 states that player i plays A exactly when he p^*-believes E and plays B exactly when he $(1-p^*)$-believes F. This is what σ^* specifies, so σ^* is a Bayesian Nash equilibrium. □

With our results behind us, we can discuss some implications. First, while it is not necessary that the distribution of ω and that of $s_i \mid \omega$ are both uniform, the result certainly does not hold for all distributions. In particular, the result holds when either the true distribution is close to uniform or the noise is small. Thus, it is only useful for the British government to avoid introducing new taxes if the prior beliefs on the true tax rate is close to uniform in a sufficiently large range. Our results also do not hold if the value of the taxes have too large an effect on the players' payoffs (i.e. if the payoff parameters change significantly with respect to ω).

6.2 The Emperor's Clothes

THE STORY. Suppose that John Doe is on his way to being the next game theorist superstar. He finally comes out with his first paper, and superficially it is a spectacular paper. However, the paper offers no real insight, a fact that John attempts to hide with mathematical complexity. And this is fairly clear to nearly everyone in the field. Nevertheless, editors start requesting the paper, departments start offering him positions, conferences start asking him to give the keynote. Why?

Presumably, no one wants to be the lone person in the field who disrespects the superstar. For example, nobody wants to be the only person *not* to invite John to a conference or a special journal issue; he might end up with a powerful enemy, even if John's research is not good. However, if everyone in the field disrespects John Doe, then everyone benefits from doing likewise, since no one wants his keynote speaker to be unpopular or his new recruit never to be invited to conferences. Thus, we can model this as a quorum game where A is the act of showing John Doe respect (e.g. inviting him to a conference), and B is an act of disrespect.[3]

We make the assumption that if in fact John Doe's research were as great as people expected, then everyone would treat him with respect. Furthermore, we assume that if a person can detect that John's research is bad, he can only approximately estimate how many others can detect this as well. We will show that, under mild conditions, if John's research is bad, no matter what fraction of people in the field can detect that his research has no insight, he will still be treated with respect. However, if people know exactly what fraction of the field know that John's research is bad, and that fraction is sufficiently high, then it is possible to treat John with disrespect. This is in stark contrast with the case where the error in a person's estimate is arbitrarily small.

THE MODEL. We model the information structure as follows: we assume that if John's research is in fact bad, then $1-\epsilon$ of the population can detect that it is bad. Everyone who can detect that it is bad has some impression of how

[3]Note that "The Emperor's New Clothes" can be seen as a metaphor for this story. John Doe is analogous to the Emperor and his colleagues are analogous to the citizens who do not, initially, publicly disrespect the obviously flawed superstar.

easy it is for others to detect how bad it is; namely, they each get a signal θ_i which is independently drawn from $\mathcal{U}[\epsilon-\delta, \epsilon+\delta]$. After observing his private signal, but not ϵ, player i can choose to play A or B.

As in the general setup, the payoff from each action is a function of the fraction of players who play B. Let r denote the fraction of players who play B. The payoffs are as follows.

$$u_i(A,r) = \begin{cases} a & : r \leq \bar{r} \\ b & : r > \bar{r} \end{cases} \qquad u_i(B,r) = \begin{cases} c & : r \leq \bar{r} \\ d & : r > \bar{r} \end{cases}$$

Theorem 5. *Let $\epsilon \sim \mathcal{U}[0,1]$ and $\theta_i \sim^{iid} \mathcal{U}[\epsilon-\delta, \epsilon+\delta]$ for all i and for some $\delta > 0$. Let σ^* be a strategy profile such that $\sigma_i^*(\theta_i) = B$ when $\theta_i \leq \bar{\epsilon}$ and $\sigma_i^*(\theta_i) = A$ when $\theta_i > \bar{\epsilon}$ for some $\bar{\epsilon} \in [\delta, 1-\delta]$. Then for $\delta \to 0$, σ^* is not a Bayesian Nash equilibrium if $p^* < \bar{r}$.*

Proof. We will use corollary 3 to show necessary conditions for σ^* to exist. Let E be the set of states in which at least $(1-\bar{r})$ fraction of the players choose A and let $F = E^c$, i.e. the set of states in which at least $\bar{r}$ fraction of the players choose B. By corollary 3, every state in E must be p^*-believed by at least $(1-\bar{r})$ fraction of the players, and every state in F must be $(1-p^*)$-believed by at least $\bar{r}$ fraction of the players.

First, observe that if the true state is ϵ and $\epsilon \in [\bar{\epsilon}-\delta, \bar{\epsilon}+\delta]$, then the fraction of players who receive a signal less than $\bar{\epsilon}$ (and thus play B) is $(1-\epsilon)\frac{\bar{\epsilon}-\epsilon+\delta}{2\delta}$. For $\epsilon > \bar{\epsilon}+\delta$, the fraction of players who play B is 0, and for $\epsilon < \bar{\epsilon}-\delta$, the fraction of players who play B is $(1-\epsilon)$. Therefore, there exists some threshold ϵ^* at which exactly $\bar{r}$ players play B, below which at least $\bar{r}$ players play B, and above which at least $(1-\bar{r})$ play A. Since $(1-\epsilon^*)\frac{\bar{\epsilon}-\epsilon^*+\delta}{2\delta}$ is the fraction of players who play B when the true state is ϵ^* and exactly $\bar{r}$ fraction of the players play B when ϵ^* is the true state, ϵ^* must satisfy the condition $(1-\epsilon^*)\frac{\bar{\epsilon}-\epsilon^*+\delta}{2\delta} = \bar{r}$. Thus, we obtain [4]

$$\epsilon^* = \frac{1+\bar{\epsilon}+\delta - \sqrt{(1+\bar{\epsilon}+\delta)^2 - 4(\bar{\epsilon}+\delta-2\delta\bar{r})}}{2} \tag{6}$$

Thus, we have that $E = \{\epsilon \mid \epsilon > \epsilon^*\}$, and $F = \{\epsilon \mid \epsilon \leq \epsilon^*\}$.

It must be the case that for any $\epsilon \in E$, at least $(1-\bar{r})$ fraction of the players must p^*-believe E. Observe that if a player receives no signal, then he believes with probability $(1-\epsilon^*)$ that the true state is in E. If the true state $\epsilon = 1 \in E$, then no players receive a signal. Therefore, if $1-\epsilon^* < p^*$, then E cannot be $(1-\bar{r})$-evident p^*-belief. For this reason, we can assume that $1-\epsilon^* \geq p^*$.

Now we will consider the necessary condition on F. It must be the case that for any $\epsilon \in F$, at least $\bar{r}$ fraction of the players must $(1-p^*)$-believe F. Since we assume $1-\epsilon^* \geq p^*$, if a player does not receive a signal, he does not $(1-p^*)$-believe F. Thus, we only consider the case in which a player receives a signal. If player i receives signal $\theta_i \in [\epsilon^*-\delta, \epsilon^*-\delta]$, then the probability that the true state is less or equal to ϵ^* (true state is in F) is $\frac{\epsilon^*-\theta_i+\delta}{2\delta}$. Therefore, player i $(1-p^*)$-believes F if and only if $\frac{\epsilon^*-\theta_i+\delta}{2\delta} \geq 1-p^*$. From this condition we obtain

$$\theta_i \leq 2\delta p^* - \delta + \epsilon^* \tag{7}$$

The fraction of players who receive such a signal is decreasing in the true state ϵ. Therefore, if at least $\bar{r}$ receive a signal satisfying equation (7) at ϵ^*, then at least $\bar{r}$ receive a signal satisfying equation (7) at all states in F. At ϵ^*, the fraction of players who receive a signal satisfying equation (7) is given by

$$(1-\epsilon^*)\frac{(2\delta-\delta+\epsilon^*)-\epsilon^*+\delta}{2\delta} = (1-\epsilon^*)p^* \tag{8}$$

This means that at least $\bar{r}$ fraction of the players will $(1-p^*)$-believe F at every state in F if and only if $p^* \geq \frac{\bar{r}}{1-\epsilon^*}$. Using equation (6), we obtain the following necessary condition:

$$p^* \geq \frac{2\bar{r}}{1-\bar{\epsilon}-\delta+\sqrt{(1+\bar{\epsilon}+\delta)^2 - 4(\bar{\epsilon}+\delta-2\delta\bar{r})}} \tag{9}$$

Equation (9) is decreasing in $\bar{\epsilon}$, so as $\delta \to 0$, σ^* is not a Bayesian Nash equilibrium if $p^* < \bar{r}$. □

We contrast this result with the scenario in which the exact value of ϵ is observed by those who can detect that John's research is bad (i.e. $\theta_i = \epsilon$). The following claim can be easily established.

Claim 3. *The strategy profile σ^* is a Bayesian Nash equilibrium if $\sigma_i^*(\theta_i) = A$ if $\epsilon \leq 1-\bar{r}$ and $\sigma_i^*(\theta_i) = B$ otherwise.*

[4]There are two roots for ϵ^*, but only one falls in the interval [0, 1]. This root is given by equation (6).

7 Discussion

We are now in a position to address some of the questions posed in the introduction, which we claimed require a formalism. Suppose that most players behave a certain way (e.g. do not revolt) in the "usual" state of the world. What conditions must hold in order for most players to behave abnormally (e.g. revolt)? Our framework demonstrates that when sufficiently many players behaves abnormally it must be commonly believed among them that it is not the "usual" state of the world. What exactly will happen in the absence of common knowledge? While mis-coordination seems to be the natural answer, our formalism demonstrates that this is not the case. Instead, a lack of common knowledge prevents people from taking an action that is different from the "usual" behavior. Is common knowledge exactly what is needed? In fact, common knowledge is not necessary. Instead, it is only necessary that the players *believe* it is likely that circumstances are not as usual, and believe it is likely that enough players believe it is likely that the circumstances are not as usual, and so on – where "enough" and "likely" can be precisely stated as a function of the payoffs in the game.

Not only can our framework abstractly answer the aforementioned questions, but it can also explain the puzzles described throughout the paper. We explained how eye-contact relates to common knowledge and coordination and prescribed that avoiding eye contact, when one is "caught" violating a social norm, can avoid undesirable social consequences. We observed that in global-like games, in which players observe a noisy signal of the true state, players cannot coordinate on continuous signals but can coordinate on discrete ones. As discussed in section 6.1, this offers an explanation for the Boston Tea Party. Furthermore, we considered situations modeling examples such as The Emperor's New Clothes, academic reputations, and the Tunisian revolution.

We believe that our framework can be broadly applied. We now informally discuss a few additional example. We begin with innuendos for which we will provide some prescriptions. Are innuendos necessary when making illicit requests? Our framework can be used to infer that a person wishing to make an illicit request without harming a relationship has to avoid explicit speech and use an innuendo. If he cannot think of an innuendo, an aptly placed cough that creates some noise, but not so much that the recipient is unlikely to hear the illicit request, would *not* do the trick. To understand why, observe that both players know that a cough laden request is likely to be heard, thereby creating common p-belief, provided "likely" is sufficiently high.

When would we wish to use an innuendo and when would we wish to go out of our way to explicitly state the obvious? If the expected behavior, in the absence of any illicit request is desirable (e.g. the relationship is healthy), an innuendo is preferred to explicit speech. On the other hand, if the expected behavior in the absence of any unusual circumstances is undesirable (e.g. two individuals are not committed to their relationship) then an explicit statement of something obvious can be useful. For example, a meaningless and costless apology, or perhaps an inconsequential bodily gesture, such as a handshake can cause two individuals to both move past a prior disagreement. The first time the words "I love you" are uttered can be quite powerful, even if both have expressed strong affection before. Such acts create common knowledge that the two are in love or "in a relationship" enabling them to coordinate on, for instance, mutual monogamy.

In some circumstances, the two players might disagree on whether the expected equilibrium is preferred. For instance, the U.S. might appreciate using Pakistan's airspace without her permission. If both expect the US to behave this way, it would be foolish for Pakistan to hawkishly[5] shoot down an American plane flying overhead. And it would be foolish of America not to take advantage of Pakistan's dovish response. Pakistan, naturally, would prefer an equilibrium in which they were expected to play hawkishly, and the US dovishly, thereby restoring Pakistan's sovereignty. A simple apology from the US could be sufficient to change the equilibria, even though the apology is mere words with no legal ramifications. Nevertheless, Pakistan was willing to expend millions of dollars in aid in order to pressure the US into apologizing, and the US was willing to sacrifice the ability to send troops through Pakistani territory in order to avoid such an apology.

Finally we discuss a couple of examples from the biological world. There are many examples of "quorum sensing," particularly among bacteria [24]. For example, bacteria emit and also detect the local concentration of certain signaling molecules. When they detect sufficiently high concentrations they produce light. Since the light is only useful when produced in sufficient quantity, the production of light can be modeled as a quorum game. Our results from section 6.2 apply to this scenario. Another example is mast seeding. Many plants benefit from synchronizing their pollen dispersal [16, 17, 27]. Since pollen predators satiate, one plant benefits from dispersing pollen when he expects many others to disperse their pollen. Such plants are known to use public cues to synchronize their pollen dispersal. Our results from

[5] The *hawk* and *dove* game is a model of two individuals fighting over a contested resource. It can framed as a asymmetric coordination game by permuting the columns of the payoff matrix.

section 6.1 can be applied to this example.

References

[1] S. Aaronson. The complexity of agreement. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, STOC '05, pages 634–643, New York, NY, USA, 2005. ACM.

[2] H. C. Andersen. *The Emperor's New Clothes*. Fairy Tales Told for Children - First Collection, 1837.

[3] R. J. Aumann. Agreeing to Disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.

[4] K. Binmore. *Game theory - a very short introduction*. Clarendon Press, 2007.

[5] H. Carlsson and E. van Damme. Global Games and Equilibrium Selection. *Econometrica*, 61(5):989–1018, 1993.

[6] B. Carp. *Defiance of the Patriots: The Boston Tea Party & the Making of America*. Yale University Press, 2010.

[7] M. S.-Y. Chwe. *Rational Ritual: Culture, Coordination, and Common Knowledge*. Princeton University Press, 2003.

[8] C. Dwork and Y. Moses. Knowledge and common knowledge in a byzantine environment: Crash failures. *Information and Computation*, 88(2):156 – 186, 1990.

[9] R. Fagin and J. Y. Halpern. Reasoning about knowledge and probability. *J. ACM*, 41(2):340–367, Mar. 1994.

[10] M. F. Friedell. On the structure of shared awareness. *Behavioral Science*, 14(1):28–39, 1969.

[11] J. Geanakoplos. Common knowledge. *Journal of Economic Perspectives*, 6(4):53–82, 1992.

[12] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *J. ACM*, 37(3):549–587, 1990.

[13] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artif. Intell.*, 54(3):319–379, Apr. 1992.

[14] J. Y. Halpern and M. R. Tuttle. Knowledge, probability, and adversaries. *J. ACM*, 40(4):917–960, Sept. 1993.

[15] J. C. Harsanyi and R. Selten. *A General Theory of Equilibrium Selection in Games*, volume 1 of *MIT Press Books*. The MIT Press, 1988.

[16] D. Janzen. *Why Bamboos Wait So Long to Flower*. Pacific Northwest Bamboo Press, 1997.

[17] D. Kelly and V. L. Sork. Mast seeding in perennial plants: Why, How, Where? *Annual Review of Ecology and Systematics*, 33(1):427–447, Jan. 2002.

[18] D. Lewis. *Convention: A Philosophical Study*. Harvard University Press, Cambridge, 1969.

[19] D. Monderer and D. Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1(2):170–190, June 1989.

[20] S. Morris. Coordination, communication and common knowledge: A retrospective on electronic mail game. *Oxford Review of Economic Policy*, 18(4), Feb. 2002.

[21] S. Morris, H. Shin, and Y. U. C. F. for Research in Economics. *Global games: theory and applications*. Cowles Foundation discussion paper. Cowles Foundation for Research in Economics, 2000.

[22] S. Morris and H. S. Shin. Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 88(3):587–97, June 1998.

[23] S. Morris and H. S. Shin. Rethinking multiple equilibria in macroeconomic modeling. In *NBER Macroeconomics Annual 2000, Volume 15*, NBER Chapters, pages 139–182. National Bureau of Economic Research, Inc, June 2001.

[24] K. H. Nealson, T. Platt, and J. W. Hastings. Cellular control of the synthesis and activity of the bacterial luminescent system. *Journal of bacteriology*, 104(1):313–322, Oct. 1970.

[25] S. Perry and J. Manson. *Manipulative Monkeys: The Capuchins of Lomas Barbudal*. Harvard University Press, 2008.

[26] S. Pinker, M. A. Nowak, and J. J. Lee. The logic of indirect speech. *PNAS*, 105(3):833–838, January 2008.

[27] M. Rees, D. Kelly, and O. N. Bjornstad. Snow tussocks, chaos, and the evolution of mast seeding.? *The American Naturalist*, 160(1):pp. 44–59, 2002.

[28] A. Rubinstein. The electronic mail game: Strategic behavior under "almost common knowledge.". *American Economic Review*, 79(3):385–91, June 1989.

A Appendix

A.1 Proof of Claim 2

Proof. Suppose $\sigma^*(\mathcal{H}) = (A, A)$ and $\sigma^*(\omega) = (B, B)$ for some $\omega \in \Omega \backslash \{(\mathcal{R}', \mathcal{G}')\}$. Let $E = \{\mathcal{H}\}$ and $F = \{\omega\}$ Then by corollary 2, we need E^c to be common $(1 - p^*)$-belief at ω.

We will make use of the inductive characterization of a common belief event provided in Lemma 1. Let us first find $B_B^{1-p}(E^c)$ and $B_D^{1-p}(E^c)$ to identify C^1. On Bob's end we have

$$\begin{aligned}
\mu(E^c|\Pi_B((\mathcal{R}', \mathcal{G})) &= \mu(E^c|\Pi_B((\mathcal{R}', \mathcal{G}')) = 1 \\
\mu(E^c|\Pi_B(\mathcal{H})) &= \mu(E^c|\Pi_B((\mathcal{R}, \mathcal{G})) = \mu(E^c|\Pi_B((\mathcal{R}, \mathcal{G}')) \\
&= \mu(\{(\mathcal{R}, \mathcal{G}), (\mathcal{R}, \mathcal{G}')\}|\{\mathcal{H}, (\mathcal{R}, \mathcal{G}), (\mathcal{R}, \mathcal{G}')\}) \\
&= \frac{p_b p_r}{p_b p_r + (1 - p_b)} < 1 - p^*
\end{aligned}$$

Hence, $B_B^{1-p^*}(E^c) = \{(\mathcal{R}', \mathcal{G}), (\mathcal{R}', \mathcal{G}')\}$.

On Dave's end we have

$$\begin{aligned}
\mu(E^c|\Pi_D(\mathcal{H})) &= 0 \\
\mu(E^c|\Pi_D((\mathcal{R}, \mathcal{G})) &= \mu(E^c|\Pi_D((\mathcal{R}', \mathcal{G})) = 1 \\
\mu(E^c|\Pi_D((\mathcal{R}, \mathcal{G}')) &= \mu(E^c|\Pi_D((\mathcal{R}', \mathcal{G}')) = 1
\end{aligned}$$

Hence, $B_D^{1-p^*}(E^c) = \{(\mathcal{R}, \mathcal{G}), (\mathcal{R}', \mathcal{G}), (\mathcal{R}, \mathcal{G}'), (\mathcal{R}', \mathcal{G}')\}$.

Therefore,

$$C^1 = B_B^{1-p^*}(E^c) \cap B_D^{1-p^*}(E^c) = \{(\mathcal{R}', \mathcal{G}), (\mathcal{R}', \mathcal{G}')\}.$$

Next we find $B_B^{1-p^*}(C^1)$ and $B_D^{1-p^*}(C^1)$ to identify C^2. On Bob's end we have

$$\begin{aligned}
\mu(C^1|\Pi_B((\mathcal{R}', \mathcal{G})) &= \mu(C^1|\Pi_B((\mathcal{R}', \mathcal{G}')) = 1 \\
\mu(C^1|\Pi_B(\mathcal{H})) &= \mu(C^1|\Pi_B((\mathcal{R}, \mathcal{G})) = \mu(C^1|\Pi_B((\mathcal{R}, \mathcal{G}')) = 0
\end{aligned}$$

Hence, $B_B^{1-p^*}(C^1) = \{(\mathcal{R}', \mathcal{G}), (\mathcal{R}', \mathcal{G}')\}$.

On Dave's end we have

$$\begin{aligned}
\mu(C^1|\Pi_D((\mathcal{R}', \mathcal{G}')) &= 1 \\
\mu(C^1|\Pi_D(\mathcal{H})) &= \mu(C^1|\Pi_D((\mathcal{R}, \mathcal{G}')) = 0 \\
\mu(C^1|\Pi_D((\mathcal{R}, \mathcal{G})) &= \mu(C^1|\Pi_D((\mathcal{R}', \mathcal{G})) \\
&= \mu(\{(\mathcal{R}', \mathcal{G})\}|\{(\mathcal{R}, \mathcal{G}), (\mathcal{R}', \mathcal{G})\}) = 1 - p_r < 1 - p^*
\end{aligned}$$

Hence, $B_D^{1-p^*}(C^1) = \{(\mathcal{R}', \mathcal{G}')\}$ and

$$C^2 = B_B^{1-p^*}(C^1) \cap B_D^{1-p^*}(C^1) = \{(\mathcal{R}', \mathcal{G}')\}.$$

Continuing similarly gives $E^{1-p^*}(E^c) = \{(\mathcal{R}', \mathcal{G}')\}$ which contradicts the fact that E^c is common $(1 - p^*)$-belief at $\omega \in \Omega \backslash \{(\mathcal{R}', \mathcal{G}')\}$. □

A.2 Proof of Lemma 3

Proof. We will first prove that if C is r-common p-belief at ω, then there exists an r-evident p-belief event E such that $\omega \in E$ and $E \subseteq B^{r,p}(C)$. In particular, we will let $E = \bigcap_{n\geq 1} C^n$. By definition of r-common p-belief, $\omega \in E$. Furthermore, $E = \bigcap_{n\geq 1} C^n \subseteq C^1 = B^{r,p}(C)$. Thus, it only remains to show that $E \subseteq B^{r,p}(E)$.

Suppose $\omega \notin B^{r,p}(E) = B^{r,p}(\bigcap_{n\geq 1} C^n)$. Then there must be strictly more than $(1-r)$ fraction of the players who do not p-believe $\bigcap_{n\geq 1} C^n$ at ω. For each player j, let m_j be the lowest integer such that $\omega \notin B_j^p(\bigcap_{n\geq 1}^{m_j} C^n)$. Now let k be the largest such integer amongst all players, i.e. $k = \max_j (m_j)$. For all j, $\bigcap_{n\geq 1}^{k} C^n \subseteq \bigcap_{n\geq 1}^{m_j} C^n$. Thus, by monotonicity of p-belief, $B_j^p(\bigcap_{n\geq 1}^{k} C^n) \subseteq B_j^p(\bigcap_{n\geq 1}^{m_j} C^n)$ for all j. Therefore, $\omega \notin B_j^p(\bigcap_{n\geq 1}^{k} C^n)$ for all j. This means that there are strictly more than $(1-r)$ fraction of the players who do not p-believe $\bigcap_{n\geq 1}^{k} C^n$ at ω, so $\omega \notin \bigcap_{n\geq 1} C^n = E$.

Now we will show that if there exists an r-evident p-belief event E such that $\omega \in E \subseteq B^{r,p}(C)$ and $E \subseteq B^{r,p}(E)$, then C is r-common p-belief at all $\omega \in E$. It suffices to show that $E \subseteq \bigcap_{n\geq 1} C^n$. We can do so by induction on n. The base case follows from the hypothesis that $E \subseteq B^{r,p}(C)$. For the inductive step, suppose that $E \subseteq C^i$ for all $i < n$. Using lemma 2 and the inductive hypothesis, we have $E \subseteq B^{r,p}(E) \subseteq B^{r,p}(\bigcap_{i<n} C^i) = C^n$. □

A.3 Proof of Theorem 2

Proof. First, we will show that if there exists a state-dependent Bayesian Nash equilibrium σ^* then there exists a non-empty $(1-\bar{r})$-evident p^*-belief event E and a non-empty s-evident $(1-p^*)$-belief event F such that $E \cap F = \emptyset$ for some $s > \bar{r}$.

Let

$$E = \{\omega \ : \ |\{i \mid \sigma_i^*(\omega) = A\}| \geq (1-\bar{r})n\},$$

i.e. the set of states at which at least $(1-\bar{r})n$ players choose A. We need to show that E is non-empty. This follows from the fact that σ^* is state-dependent; if a player plays A on some state, then he must believe that there is some state in which at least $(1-\bar{r})$ fraction of the players choose A. If this were not true, then he would best respond by playing B.

Now suppose for contradiction that there exists some $\hat{\omega} \in E$ such that $\omega \notin B^{1-\bar{r},p^*}(E)$. By construction of E, at least $(1-\bar{r})$ fraction of the players play A at $\hat{\omega}$. This implies that there exists at least one player i who plays A on $\hat{\omega}$ but does not p^*-believe E. By construction of E, strictly less than $(1-\bar{r})$ play A outside of E. Therefore, i believes with probability strictly less than p^* that he will get payoff a from playing A and c from playing B at $\hat{\omega}$.[6] More precisely,

$$u_i(A|\Pi_i(\hat{\omega})) - u_i(B|\Pi_i(\hat{\omega})) < p^*(a-c) + (1-p^*)(b-d) = 0$$

Therefore, σ^* is not a Bayesian Nash equilibrium. An analogous proof shows that $F \subseteq B^{s,1-p^*}(F)$, where $F = \{\omega \mid |\{i \mid \sigma_i^*(\omega) = B\}| \geq sn\}$ for some $s > \bar{r}$.[7]

Now we will show that if there exists a non-empty $(1-\bar{r})$-evident p^*-belief event E and a non-empty s-evident $(1-p^*)$-belief event F such that $E \cap F = \emptyset$ and $s > \bar{r}$, then there exists a state-dependent Bayesian Nash equilibrium σ^*. We will do so by constructing such a Bayesian-Nash equilibrium.

First, we will recursively define the sequence of events E^t and F^t as follows. Let $E^0 = E$ and $F^0 = F$. For all $t \geq 1$, let $E^t = B^{1-\bar{r},p^*}(E^{t-1})$ and $F^t = \bigcup_{q<p^*} B^{s,1-q}(F^{t-1})$. Finally, let $\hat{E} = \bigcup_{t=0}^{\infty} E^t$ and $\hat{F} = \bigcup_{t=0}^{\infty} F^t$.

Now let σ^* be a strategy profile such that

$$\begin{aligned}
\sigma_i^*(\omega) &= A & &\forall \omega \in B_i^{p^*}(\hat{E}) \\
\sigma_i^*(\omega) &= B & &\forall \omega \in B_i^{1-p^*}(\hat{F}) \\
\sigma_i^*(\omega) &= B & &\forall \omega \in B_i^{p^*}(\hat{E})^c \cap B_i^{1-p^*}(\hat{F})^c
\end{aligned}$$

We will show that σ^* is a Bayesian Nash equilibrium using the following 4 steps:

1. $E^t \subseteq E^{t+1}$ and $F^t \subseteq F^{t+1}$.

2. $E^t \cap F^t = \emptyset$ for all $t \geq 0$.

[6] As noted in the setup, we can ignore the contribution of player i to the fraction of players who play B because n is sufficiently large.

[7] We use strict inequalities to construct F because strictly more than $\bar{r}$ players must play B in order to achieve payoffs b or d.

3. For all $\omega \in B_i^{p^*}(\hat{E})$, player i will be no better off by playing B.

4. For all $\omega \in B_i^{p^*}(\hat{E})^c \cap B_i^{1-p^*}(\hat{F})^c$, player i will be no better off by playing A.

The first two steps show that σ^* is well defined. The final two steps, combined with an analogous argument for all $\omega \in \hat{F}$ show that σ^* is a Bayesian Nash equilibrium.

Proof of step 1: We will show that $E^t \subseteq E^{t+1}$ by induction on n. For the base case, note that $E^0 = E \subseteq B^{1-\bar{r},p^*}(E) = E^1$ because E is a $(1-\bar{r})$-evident p^*-belief event. Now suppose $E^{t-1} \subseteq E^t$. Then by lemma 2 (monotonicity of r,p-beliefs), $E^t = B^{1-\bar{r},p^*}(E^{t-1}) \subseteq B^{1-\bar{r},p^*}(E^t) = E^{t+1}$. An analogous argument proves that $F^t \subseteq F^{t+1}$.

Proof of step 2: We will show that $E^t \cap F^t = \emptyset$ for all $t \geq 0$ by induction on t. For the base case, note that $E^0 \cap F^0 = E \cap F = \emptyset$. Now suppose $E^t \cap F^t = \emptyset$. Then $\mu(E^t \mid \Pi_i(\omega)) + \mu(F^t \mid \Pi_i(\omega)) \leq 1$ for all ω and for all i. If $\omega \in E^{t+1}$, then by definition of E^{t+1}, for at least $(1-\bar{r})$ fraction of the players, $\mu(E^t \mid \Pi_i(\omega)) \geq p^*$. This implies that for at least $(1-\bar{r})$ fraction of the players, $\mu(F^t \mid \Pi_i(\omega)) \leq 1-p^*$. Therefore, for at most $\bar{r}$ fraction of the players, $\mu(F^t \mid \Pi_i(\omega)) \geq 1-q$. Since $s > \bar{r}$, $\omega \notin F_i^{t+1}$.

Proof of step 3: Suppose $\omega \in B_i^{p^*}(\hat{E})$. Then from the definition of $\hat{E}$, $\mu(E^t \mid \Pi_i(\omega)) \geq p^*$ for some $t \geq 0$. First, consider the case in which $t = 0$, i.e. $\mu(E \mid \Pi_i(\omega)) \geq p^*$. Since $E \subseteq B^{1-\bar{r},p^*}(E)$ and $E \subseteq \hat{E}$, by monotonicity of r,p-beliefs, $E \subseteq B^{1-\bar{r},p^*}(\hat{E})$. Therefore, at ω, player i believes with probability at least p^* that at least $(1-\bar{r})$ players will play A. Thus, player i will be no better off by playing B.

Now consider the case in which $t > 0$. Since $E^{t-1} \subseteq \hat{E}$, by monotonicity of r,p-beliefs, $B^{1-\bar{r},p^*}(E^{t-1}) \subseteq B^{1-\bar{r},p^*}(\hat{E})$. Therefore, $\mu(B^{1-\bar{r},p^*}(\hat{E}) \mid \Pi_i(\omega)) \geq p^*$, and hence, at ω, player i believes with probability at least p^* that at least $(1-\bar{r})$ players will play A. Thus, player i will be no better off by playing B.

Proof of step 4: Suppose $\omega \in B_i^{p^*}(\hat{E})^c \cap B_i^{1-p^*}(\hat{F})^c$. Suppose $\omega \in B_i^{p^*}(B^{1-\bar{r},p^*}(\hat{E}))$. Then player i believes with probability at least p^* that at least $(1-\bar{r})$ fraction of the players p^*-believe $\hat{E}$. Since $E^t \subseteq E^{t+1}$ for all t, there must be some t^* such that $\omega \in B_i^{p^*}(B^{1-\bar{r},p^*}(E^{t^*}))$. Since $B^{1-\bar{r},p^*}(E^{t^*}) = E^{t^*+1} \subseteq \hat{E}$, $\omega \in B_i^{p^*}(\hat{E})$. This is a contradiction, so $\omega \notin B_i^{p^*}(B^{1-\bar{r},p^*}(\hat{E}))$. Therefore, at ω, player i believes with probability strictly less than p^* that at least $(1-\bar{r})$ fraction of the players will play A. This implies that player i would be no better off by playing A at ω. □