



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Communication Equilibria and Bounded Rationality


Nikhil Vellodi

No 955

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK



Communication Equilibria and Bounded Rationality

Nikhil Vellodi*

University of Warwick

September 17, 2010

Abstract

In this paper, we generalize the notion of a communication equilibrium (Forges 1986, Myerson 1986) of a game with incomplete information by introducing two new types of correlation device, namely *extended* and *Bayesian* devices. These new devices explicitly model the ‘thinking process’ of the device, i.e. the manner in which it generates outputs conditional on inputs. We proceed to endow these devices with both information processing errors, in the form of non-partitional information, and multiple transition and prior distributions, and prove that these two properties are equivalent in this context, thereby generalizing the result of Brandenburger, Dekel and Geanakoplos (1988). We proceed to discuss the Revelation Principle for each device, and conclude by nesting a certain class of ‘cheap-talk’ equilibria of the underlying game within Bayesian communication equilibria. These so-called *fallible talk* equilibria cannot be generated by standard communication equilibria.

*The author is available at n.vellodi@gmail.com. I am greatly indebted to Françoise Forges and Indra Ray for their numerous suggestions, Sushama Murty for her initial guidance, Peter Hammond for invaluable discussions, and especially my supervisor Herakles Polemarchakis, for his patience, good humour and generosity. Naturally, all errors and omissions are the author’s sole responsibility.

Contents

1	Introduction	3
2	Background	4
2.1	Inconsistent Beliefs and Non-Partitional Information	4
2.2	Correlated and Communication Equilibria	6
2.3	Canonical Equilibria and the Revelation Principle	8
2.3.1	Generalized and Subjective Correlated Equilibria	10
2.4	Cheap Talk Games and Communication Equilibria	10
3	Extended Communication Equilibria	12
3.1	Potential Drawbacks	20
4	Bayesian Communication Equilibria	21
5	Further Topics	30
5.1	The Revelation Principle and Generalized Bayesian Devices	30
5.2	Fallible Talk and Bayesian Communication Equilibria	33
6	Conclusion	35
A	Appendix A	39
B	Appendix B	40
C	Appendix C	44

1 Introduction

Much has been written on the topics of the Common Prior assumption, non-partitional information and communication equilibria. To the author’s best knowledge, nothing has been written on the intersection of all three areas. This is the broad objective of the present paper.

It is well known that if players are fully rational and share a common prior on the uncertain states of the world, then there is no incentive in equilibrium to bet whilst trading (see [Milgrom, Stokey, 1982]). The inverse is often observed in reality; people who share different probability assessments regarding an event often bet against each other, thus expanding the set of equilibrium payoffs available. This behaviour is rational in the sense that betting can form part of a Nash equilibrium. However, there are arguments against admitting so-called ‘heterogenous priors’; see [Harsanyi, 1967] and [Aumann, 1987]. It is also well known that information processing that admits bounded rationality generates similar phenomena (see [Geanakoplos, 1989]). The paper of Brandenburger, Dekel and Geanakoplos (1989) represents a concise attempt to link heterogenous priors and bounded rationality, within the context of correlated equilibria. In this paper, we offer a natural extension of their work to communication equilibria. It is easy to think of situations in which communication devices, which essentially act as a third-party arbitrators between players in Bayesian games, could exhibit some form of bounded rationality or imperfection. Take a market mechanism in a production economy under uncertainty, which receives information regarding players’ endowments and preferences, and issues advice regarding prices. It is certainly conceivable that such a mechanism could operate under information processing imperfections. However, in order to study such situations, we must first explicitly model the manner in which the device interprets information, i.e. its ‘thinking process’. In this paper, we construct two new types of communication device that allow us to do precisely this.

We begin by giving a brief overview of the existing work on inconsistent beliefs and non-partitional information, communication equilibria, canonical equilibria and the Revelation Principle and cheap talk extensions. Chapter 2 introduces the concept of an *extended communication device*, or simply extended devices. Such devices consist of an input space, and output space and a transition probability from a subset of the power set of the input space to the output space. We provide examples that demonstrate how such devices offer Pareto-improving outcomes to standard devices, through a process akin to ‘betting’. We then introduce the notion of *decision-theoretic equivalence*, and show that the set of generalized extended communication equilibria is equal to the set of subjective extended communication

equilibria, via this equivalence relation. Chapter 3 introduces the notion of a *Bayesian communication device*, or simply Bayesian devices. We show that extended devices are essentially a reduced form version of Bayesian devices. We redefine the notion of decision-theoretic equivalence so that it retains tractability in this new setting, and prove that the set of generalized Bayesian devices is equal to the set of subjective Bayesian devices. Chapter 4 discusses briefly the problem of classifying the set of canonical extended and Bayesian equilibria, or equivalently proving a revelation principle in this case. We show how the problem is trivial for subjective Bayesian devices, but that the problem is more complex for generalized Bayesian devices. Chapter 5 discusses how Bayesian communication equilibria nest a certain class of cheap talk equilibria of the game, in which players can communicate directly, but such communication is prone to error. As such, this process is termed *fallible talk*. We conclude by discussing possible further extensions.

2 Background

2.1 Inconsistent Beliefs and Non-Partitional Information

Consider the following situation. A decision-maker is faced with uncertainty, represented by a state-space Ω . He already possesses beliefs about these uncertain states, represented by a *prior* probability distribution π over Ω . When a state is realized, he can perceive it up to some level of accuracy. This is represented by his *information function*:

$$\mathcal{P} : \Omega \longrightarrow 2^\Omega,$$

where 2^Ω is the power set of Ω . Each element of the image space $\mathcal{P}(\Omega)$ is referred to as an *information set*, and we say that he ‘knows $\mathcal{P}(\omega)$ at ω ’.

Having perceived the information set the true state lies in, he now updates his prior distribution to form his *posterior* distribution, accounting for this new information through a process of *Bayesian updating*:

$$\bar{\pi}(\omega'|\mathcal{P}(\omega)) = \begin{cases} \frac{\pi(\omega')}{\sum_{\hat{\omega} \in \mathcal{P}(\omega)} \pi(\hat{\omega})} & \text{if } \omega' \in \mathcal{P}(\omega) \\ 0 & \omega' \notin \mathcal{P}(\omega) \end{cases}$$

For example, suppose $\Omega = \{1, 2, 3\}$, $\pi(1) = \pi(2) = \pi(3) = \frac{1}{3}$, and $\mathcal{P}(1) = \mathcal{P}(2) = \{1, 2\}$, $\mathcal{P}(3) = \{3\}$.

Now suppose the true state is 1. Then the posteriors become $\bar{\pi}(1) = \frac{1}{2}, \bar{\pi}(2) = \frac{1}{2}, \bar{\pi}(3) = 0$.

This example uses a special form of information function, in that the image under \mathcal{P} forms a partition of Ω . In this case, we say that the information function induces an *information partition* of Ω ; more specifically, if

1. $\omega \in \mathcal{P}(\omega), \forall \omega \in \Omega$ (*Non-Delusion*)
2. $\omega' \in \mathcal{P}(\omega) \Rightarrow \mathcal{P}(\omega') = \mathcal{P}(\omega)$

Within the context of Bayesian games, such information structures are generally assumed to hold. Indeed, the ‘type-space’ representation of uncertainty implicitly assumes partitional information. The standard form of a Bayesian game assumes a state space Ω , a type space T_i for each i , and a surjective signal function $\tau_i : \Omega \rightarrow T_i$ (see [Osborne, Rubinstein, 1994]). Then if $\tau_i(\omega) = t_i$, we have that $\mathcal{P}_i(\omega) = \tau_i^{-1}(t_i)$, where \mathcal{P}_i is player i ’s information function. That τ_i is a function implies that \mathcal{P}_i is partitional, in particular, the partition \mathcal{P}_i is generated by the equivalence relation

$$\omega \sim_i \omega' \Leftrightarrow \tau_i(\omega) = \tau_i(\omega')$$

If the information function does not necessarily induce a partition, we call it a *possibility correspondence* (see [Geanakoplos, 1989]). It can be shown that such structures embody a departure from full rationality, or at the very least, error-free information processing. In a possibility correspondence, the existence of ‘overlaps’ implies that agents do not necessarily ‘know what they know’, i.e. it may not be the case that $\omega' \in \mathcal{P}(\omega) \Rightarrow \mathcal{P}(\omega') \subset \mathcal{P}(\omega)$. There are many forms of such ‘bounded rationality’ that have been studied extensively. See for example [Geanakoplos, 1989], [Samet, 1990] and [Shin, 1993].

It is clear how the scenario above can be extended to an arbitrary n -person set-up. In this case, the assumption that the prior distribution is identical for all players is precisely the Common Prior assumption (CPA), or Harsanyi doctrine.¹ Much has been written on the role of the CPA in Bayesian games.² Indeed, that posterior distributions are allowed to differ on account of diverging information represents faithfully Harsanyi’s (1967) original rationale, which argues that differences in player’s beliefs should stem purely from differences in information, i.e. that “...there is no rational basis for people who have always been fed precisely the same information to maintain different subjective probabilities”

¹Departure from the CPA is variously referred to as subjective/different/heterogenous/inconsistent priors/beliefs.

²See [Aumann, 1987].

([Aumann, 1976]). As such, subjective priors could be explained by errors in information processing. This is one of the central propositions of the present paper. There are however several defenses of inconsistent beliefs. See, for example, [Morris, 1995]. Morris cites various examples where learning opportunities may be limited, thus resulting in different priors. Indeed, such an observation implicitly assumes that the formation of the prior is far from arbitrary, and moreover that a common prior assumes boundless opportunities to learn prior to the commencement of the game. The present question is: how can inconsistent beliefs be justified within the context of games with communication? Far from being an advocacy of heterogenous priors, this paper suggests that heterogenous priors can be explained via bounded rationality, specifically non-partitional information structures.

Brandenburger, and Dekel and Geanakoplos 1988 - henceforth BDG - show that in the context of correlated equilibrium (defined below), heterogenous priors and non-partitional information essentially amount to the same thing. We will discuss this result at length below.

2.2 Correlated and Communication Equilibria

The concept of a correlated equilibrium was proposed in [Aumann, 1967] for normal-form games. It poses both the first credible alternative to (mixed) Nash equilibrium and the first real attempt at incorporating communication into a game theoretic solution concept. Essentially, the idea behind correlated equilibrium is to introduce a layer of uncertainty into the solution in order to help correlate players' strategies. Take the familiar example of a crossroads: two cars approach a crossroads at right-angles. Their pure strategies are 'stop' and 'go'. Clearly, for some natural choice of payoffs, there exist Nash equilibria that impose inequitable delays, or may even result in an accident. However, installation of a traffic light, with two states corresponding to 'stop' and 'go', can clearly help correlate player's actions to achieve a pareto-superior outcome to any Nash equilibrium.³ This example demonstrates how the uncertainty in a correlated equilibrium can be represented by a 'device' - in this case, a traffic light - which generates outputs that the players use to make their choices. With this in mind, we present a formal definition of correlated equilibrium, based on [Forges, 2009].

Definition 1. A *correlation device* is a tuple $d = (\Omega, (\mathcal{P}_i)_{i \in N}, \pi)$, where Ω is the state space, \mathcal{P}_i is player i 's partitional information function and π is the prior measure over Ω .

Now let $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ be a game. Let $d = (\Omega, (\mathcal{P}_i)_{i \in N}, \pi)$ be a correlation device. Then a

³This is an example of a canonical correlated equilibrium, which will be discussed at length later.

correlated equilibrium of Γ is a tuple $(d, (f_i)_{i \in N})$, where d is a correlation device, and f is a Nash equilibrium of the game Γ^d , constructed by adding the device d to Γ .

Γ^d is played as follows: First, the state of nature is realized. Second, each player is informed of their private information, using the \mathcal{P}_i 's. Third, the player's make their moves, conditional on this information.

Specifically, the equilibrium conditions are:

$$\sum_{\omega' \in \mathcal{P}_i(\omega)} \pi(\omega' | \mathcal{P}_i(\omega)) u_i(f(\omega)) \geq \sum_{\omega' \in \mathcal{P}_i(\omega)} \pi(\omega' | \mathcal{P}_i(\omega)) u_i(a_i, f^{-i}(\omega)), \quad \forall i \in N, \forall \omega \in \Omega : \pi(\omega) > 0, \forall a_i \in A_i$$

The notion of a *communication equilibrium* was introduced in [Myerson, 1982], later formalized in [Myerson, 1986] and generalized in [Forges, 1986], as an extension of correlated equilibrium to Bayesian games. Correlated equilibrium involves unidirectional communication. This may entail a loss of efficiency in the incomplete information case, where players now have private information that could be put to use by the device in generating its output. Recall the example of the traffic light. Assume now that each car could be approaching at two speeds, ‘fast’ and ‘slow’, could be either ‘far’ or ‘near’ to the cross-roads, and that the other car is uncertain about these parameters. If the traffic light could collect information regarding each car’s speed and distance from the cross-roads, it could generate outcomes pareto-superior to those without such capabilities.

Thus, if we extend the correlation device to allow bi-directional communication, i.e. to and from the players, we can successfully utilize the information contained in players’ types to broaden the set of equilibrium payoffs. This is essentially the intuition behind communication equilibria. We formalize this in the following definition:⁴

Definition 2. An *communication device* is a tuple

$$d_c = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi)),$$

where S_i is the set of inputs from i to the device, M_i is the set of private outputs from the device to i , π is the transition probability the device uses to generate output vectors given input vectors, and \mathcal{P}_d^i is the device’s possibility correspondence on $S = \prod_{i \in N} S_i$ for each i .

⁴This definition follows Forges (1986). Later, we will show how Myerson’s definition is nested in Forges’.

Now let $\Gamma_B = (N, (T_i)_{i \in N}, p, (A_i)_{i \in N}, (u_i)_{i \in N})$ be a Bayesian game. Let

$$d_c = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi))$$

be a communication device. The game Γ_B^d formed by attaching the device d_c to Γ_B is played as follows: first, the players' types are realized according to p . Next, the players choose an input $s_i \in S_i$ to communicate to the device⁵. The device then generates a private output $m_i \in M_i$ according to the transition measure π for each player, conditional on having received the input vector $s = \prod_{i \in N} s_i$. Finally, each player chooses an action in A_i , conditional on their type, their chosen input and the device's private output. Hence, a strategy for player i is a pair (σ_i, δ_i) , where

$$\begin{aligned} \sigma_i &: T_i \longrightarrow S_i \\ \delta_i &: S_i \times M_i \times T_i \longrightarrow \Delta A_i \end{aligned}$$

In other words, σ_i chooses player i 's input, given their type, and δ_i chooses player i 's action in Γ , given their input, private output and type. A strategy profile is defined in the usual way over the product spaces.

An communication equilibrium (CE) of Γ_B is a Nash equilibrium of the game Γ_B^d , for a given d_c .

2.3 Canonical Equilibria and the Revelation Principle

The traffic light examples hint at the existence of a special class of devices, in which the output space is the action space and (in the context of communication devices) the input space is the type space. Such devices are called *canonical* devices, since the input/output spaces are the canonical choices. If we had chosen the traffic lights to emit three different coloured lights, which the drivers then had to interpret, this would no longer be a canonical device. Formally:

Definition 3. A canonical correlation device is one in which $\Omega = A$. A canonical communication device is one in which $S = T$ and $M = A$.

There are two advantages of dealing with such devices. On the technical front, the equilibrium condi-

⁵We assume no randomization is allowed, as the set S could be modified to incorporate mixing if desired.

tions, which now become:

$$\sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \sum_{a \in A} \pi(a|t) u^i(t, a) \geq \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \sum_{a \in A} \pi(a|s_i, t_{-i}) u^i(t, f(a_i), a_{-i}),$$

$$\forall i \in N, \forall t_i, s_i \in T_i, \forall f : A_i \rightarrow A_i$$

form a system of linear equations, whose solution is easily obtained.⁶ On the practical front, such a representation makes more concrete the role of the device as a third party mediator, which issues direct recommendations on actions, given information regarding players types. Diagrammatically, we have:

$$\text{Non-canonical communication device: } T \xrightarrow{\sigma} S \xrightarrow{\pi} M \xrightarrow{\delta} A$$

$$\text{Canonical communication device: } T \xrightarrow{\sigma} T \xrightarrow{\pi} A \xrightarrow{\delta} A$$

Whilst the definition of strategies remain unchanged, the definition of equilibrium needs discussion. One might be forgiven for assuming that a canonical communication equilibrium is any Nash equilibrium of the game formed by attaching a canonical device to a Bayesian game Γ_B . However, [Forges, 1986], [Myerson, 1986] and subsequent authors include the extra requirement that, in a canonical equilibrium, players must be honest and obedient, i.e. they report their true type to the device, and take the action recommended to them by the device. Indeed, the inequalities stated above are for such an equilibrium. Whilst it may seem unintuitive to allow for dishonesty and disobedience in equilibrium, it is certainly feasible. Hence, we define a canonical equilibrium to be simply an equilibrium with respect to a canonical device, and call a canonical equilibrium in which players are honest and obedient an *honest-and-obedient* canonical equilibrium.

The Revelation Principle is a statement that applies to any equilibrium concept involving communication. Informally, it states that for any non-canonical equilibrium, there is a payoff-equivalent honest-and-obedient canonical equilibrium. The construction involved in the proof is essentially the same for correlated and communication equilibria; if one treats the device as a 'black box', one can simulate the actions players would have played in the non-canonical case, and enforce the canonical device to suggest these actions as its output. Clearly, there is no incentive to deviate from such recommendations, since the original profile was also in equilibrium.

⁶This assumes that the action and type spaces are finite. We will maintain this assumption throughout this paper.

2.3.1 Generalized and Subjective Correlated Equilibria

So far, we have dealt with devices that involve consistent beliefs, whether this be a common prior or a uniform transition distribution, and partitional information. The relaxation of these restrictions was examined in BDG in the case of correlation devices. Devices which use possibility correspondences, rather than information functions are called *generalized correlation devices*, whilst devices that use different priors on Ω are called *subjective correlation devices*.⁷ Amongst the results they prove are that the set of generalized equilibria and the set of subjective equilibria are decision-theoretically equivalent, where this concept is a stronger version of payoff equivalence that requires actions to be equal on classes of information sets that are isomorphic to each other. This is an important result for two reasons. First, it offers an implicit defense of the CPA, as it shows how permitting heterogenous priors is tantamount to bounded rationality. Second, on the flip side, it shows that if bounded rationality is to be taken seriously, or at least is of interest from a modeling perspective, then so too must subjective priors. They also characterize the set of so-called *generalized correlated equilibrium distributions*, which are essentially the canonical counterpart of equilibria generated using generalized correlation devices. They prove a form of the Revelation Principle in the case of generalized correlated equilibria by characterizing the set of canonical generalized correlated equilibria. We will discuss the extension of this result to an analogously defined communication equilibrium.

It is easy to prove the Revelation Principle with respect to subjective correlated equilibrium:

Proposition 1. *Let (d, α) be a subjective correlated equilibrium for the game $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ with respect to the subjective correlation device $d = (\Omega, (\mathcal{P}_i)_{i \in N}, (\pi)_{i \in N})$. Then there exists a payoff equivalent canonical subjective correlated equilibrium in which players are obedient.*

Proof. Set $\hat{\Omega} = A$. Set $q_i(a) = \pi_i(\{\bar{\omega} \in \Omega | \alpha(\bar{\omega}) = a\})$. Let $\hat{\mathcal{P}}_i(a) = \{b \in A | b_i = a_i\}$, where a_i is the i th component of the action profile a . Finally, let $\hat{\alpha}_i(a) = a_i$. Then $d_{can} = (\hat{\Omega}, (\hat{\mathcal{P}}_i)_{i \in N}, (q_i)_{i \in N})$ is a subjective canonical device, and $(d_{can}, \hat{\alpha})$ is the desired subjective canonical correlated equilibrium. \square

2.4 Cheap Talk Games and Communication Equilibria

Whilst communication equilibria are the natural extension of correlated equilibria to games with incomplete information, their original formulation was in the context of a set of stylized principal-agent

⁷They do not explicitly use the language of devices, but the definitions are equivalent.

games, where an informed agent reports information regarding his type to a totally uninformed principal, who then takes an action that determines the payoff for both.⁸ Such games have been termed *cheap-talk* games, as the information transmitted by the agent has no direct effect on the payoff to either player. The earliest examples of the analysis of such games are to be found in [Crawford, Sobel, 1982] and [Green, Stokey, 1980].⁹ One can extend this scenario to allow for more general forms of communication, i.e. allowing both principal and agent to have private information, allowing both to have information transmission roles and actions, etc.

The similarities to the communication device setup are clear, with one important distinction. With communication devices, all information transmission is private; the device acts as a third-party mediator for communication between players, and does not reveal its received inputs or generated outputs publicly.¹⁰ Clearly, in the cheap-talk scenario, the device is totally removed from the system, and the players communicate directly. As a result, it is to be expected that cheap talk equilibria may be contained within private communication equilibria, but maybe not the converse. In other words:¹¹

Definition 4. *Let $\Gamma_B = (N, (T_i)_{i \in N}, p, (A_i)_{i \in N}, (u_i)_{i \in N})$ be a Bayesian game. Then a **cheap-talk extension** $ext(G)$ of Γ_B consists of T stages of costless, unmediated communication before G is played, in which at each $t \in T$, each i simultaneously selects a message in M_t^i and transmits it to a subset of players. After the communication phase, players choose their actions. If the communication is before players learn their types, then it is called an **ex ante** phase. If it occurs after they learn their types but before they choose their actions, then it is called an **interim** phase.*

Remark 1. *The payoffs generated by all $ext(G)$ is contained in the space of payoffs generated by all communication equilibria.*

For a formal discussion of this statement and its proof, see [Forges, 2009]. Of course, the duality represented in Remark 1 depends entirely on the manner of communication allowed for. It is natural, then, to assume that by generalizing the notion of a communication equilibria, we could find a suitably generalized notion of cheap-talk that mirrors this duality. This is the aim of Chapter 5.2, which defines a cheap-talk extension corresponding to our notion of a Bayesian communication equilibrium.

⁸See [Myerson, 1982], [Forges, 1985].

⁹The latter involves a more general setup than described, and will be discussed later.

¹⁰This idea was pursued in [Lehrer, 1996], which allows the device to issue public announcements.

¹¹This definition is taken from [Forges, 2009].

3 Extended Communication Equilibria

We are now ready to introduce the first of our generalizations of the concept of a communication device. Combining the concepts outlined above, the most general form of a communication device presently available follows the scheme:

$$T \xrightarrow{\sigma} S \xrightarrow{\pi} M \xrightarrow{\delta} A$$

We can extend this to cover the case where the device uses different transition probabilities to generate its output. Naturally, this may be termed a subjective communication device, following the scheme

$$T \xrightarrow{\sigma} S \xrightarrow{\pi_i} M \xrightarrow{\delta} A$$

In words, the device receives an input vector from the players and then uses different transition distributions to calculate private output vectors for each player. Such devices are open to attack by proponents of the CPA as discussed. We aim to analyze such devices by ‘opening up’ the manner in which the transition probabilities are calculated, and then demonstrating that such subjective devices can be viewed as having bounded rationality in the form of non-partitional information. Let us begin with a formal definition.

Definition 5. *A subjective extended communication device (SECD) is a tuple*

$$d_c = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (\mathcal{F}_i)_{i \in N}),$$

where S_i is the set of inputs for i to the device, M_i is the set of private outputs from the device to i , \mathcal{F}_i is a partition of S for each i , and π_i is the subjective transition distributions the device uses to generate output vectors given input vectors. More precisely,

$$\pi : \mathcal{F}_i \longrightarrow \Delta M$$

An *extended communication device (ECD)* is a subjective extended communication device where, for $f_i \in \mathcal{F}_i$ and $f_j \in \mathcal{F}_j$,

$$\pi_i(f_i) = \pi_j(f_j) \text{ when } f_i = f_j$$

i.e. the distributions over M conditional on common information sets agree. A *subjective extended*

communication equilibrium (SECE) of Γ_B is a Nash equilibrium of the game formed by attaching a SECD to Γ_B .

Remark 2. Take a Bayesian game $\Gamma_B = (N, (T_i)_{i \in N}, p, (A_i)_{i \in N}, (u_i)_{i \in N})$, and a SECD $d_e = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (\mathcal{F}_i)_{i \in N})$. Then playing the strategy profile (σ, δ) yields an expected utility to player i of type t_i of

$$\sum_{t \in T} p(t_{-i} | t_i) \sum_{m \in M} \sum_{a \in A} \pi_i(m | f_{i, \sigma(t)}) u_i(\delta(a | m, \sigma(t), t), \sigma, t)$$

This definition allows us to associate each $s \in S$ with an element $f_i \in \mathcal{F}_i$. Let $f_{i,s}$ denote this element. We say that the device ‘knows $f_{i,s}$ for channel i at s ’.¹²

The device operates as follows. It receives an input vector from the players. However, whereas a communication device has full certainty when reading this input vector, this new device does not, reading it to within a degree of uncertainty prescribed by the information encoded in the partitions \mathcal{F}_i . The device then generates an output conditional on this information.

Remark 3. If $\mathcal{F}_i = S$, and if $\pi_i = \pi$ for all i , then the SECD $d_e = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (\mathcal{F}_i)_{i \in N})$ is a communication device.

Indeed, we can show that the notion of a SECD does not generalize that of a subjective communication device, thereby verifying that extended devices generalize communication devices only in their operation, not in their equilibrium payoff sets.

Proposition 2. Every subjective communication device can be simulated by an SECD and vice versa.

Proof. Take a subjective communication device $d_c = (S, M, Q_i)$. Set $\mathcal{F}_i = S$ for all i . Then $f_{i,s} = s$ for all $i \in N, s \in S$, so set

$$\pi_i(m | f_{i,s}) = Q_i(m | s), \quad \forall m \in M, \forall i \in N$$

Then $d_E = (S, M, \mathcal{F}_i, \pi_i)$ is an equivalent SECD. Conversely, let $d_E = (S, M, \mathcal{F}_i, \pi_i)$ be a SECD. Let

$$Q_i(m | s) = \pi_i(m | f_{i,s}), \quad \forall m \in M, \forall s \in f_{i,s}, i \in N$$

Then $d_c = (S, M, Q_i)$ is an equivalent SCD. □

¹²Henceforth, we call each of the i processing routes within the device ‘channels’.

We give an example to demonstrate how such devices can generate equilibrium payoffs that pareto dominate those achievable through communication devices.

Example. Consider $\Gamma_B = (\{1, 2\}, (T_i), (p_i), (A_i), (u_i))$, the game of Bayesian Matching Pennies: $T_1 = \{s_1, t_1\}$, $T_2 = \{t_2\}$, $p_1(s_1) = p_1(t_1) = \frac{1}{2}$, $A_1 = \{T, B\}$, $A_2 = \{L, R\}$, and von-Neumann-Morgenstern payoffs are given in the tables below:

s_1	L	R
T	1,-1	-1,1
B	-1,1	1,-1

t_1	L	R
T	-1,1	1,-1
B	1,-1	-1,1

Define the functions $\mu_1, \mu_2 : T \rightarrow \Delta A$ as follows:¹³ If

$$\boldsymbol{\mu} = \begin{pmatrix} \mu_1((T, L)|s_1) & \mu_2((T, L)|s_1) \\ \mu_1((T, R)|s_1) & \mu_2((T, R)|s_1) \\ \mu_1((B, L)|s_1) & \mu_2((B, L)|s_1) \\ \mu_1((B, R)|s_1) & \mu_2((B, R)|s_1) \\ \mu_1((T, L)|t_1) & \mu_2((T, L)|t_1) \\ \mu_1((T, R)|t_1) & \mu_2((T, R)|t_1) \\ \mu_1((B, L)|t_1) & \mu_2((B, L)|t_1) \\ \mu_1((B, R)|t_1) & \mu_2((B, R)|t_1) \end{pmatrix}$$

Then let

¹³In the notation for μ , we assume $T \cong T_1$ by projection onto T_1 , i.e. we drop the need to include the label t_2 .

$$\boldsymbol{\mu} = \begin{pmatrix} \frac{1+\epsilon}{2} & \frac{1-\epsilon}{2} \\ 0 & 0 \\ \frac{1-\epsilon}{2} & \frac{1+\epsilon}{2} \\ 0 & 0 \\ \frac{1-\epsilon}{2} & \frac{1+\epsilon}{2} \\ 0 & 0 \\ \frac{1+\epsilon}{2} & \frac{1-\epsilon}{2} \\ 0 & 0 \end{pmatrix}$$

Clearly, $\boldsymbol{\mu}$ is a pair of transition distributions from T to A . Let $\mathcal{F}_i = T$ for all i . Then

$d_c^c = \left((T_i)_{i \in N}, (A_i)_{i \in N}, (\mu_i)_{i \in N}, (\mathcal{F}_i)_{i \in N} \right)$ is a SECD.¹⁴ Furthermore, if (α_c, δ_c) is an honest-and-obedient strategy profile, then $(d_c^c, (\sigma_c, \delta_c))$ is a SECE of Γ_B . In this equilibrium, each player receives an expected utility of $\frac{\epsilon}{2}$. It is clear that any CE can yield an expected utility of at best 0. Hence, even a small deviation in the device's beliefs over the signal space can cause Pareto-improving outcomes in equilibrium.¹⁵ For more details on this example, see Appendix A.

The causes of this payoff improvement are essentially identical to those outlined in [Aumann, 1967]; by using differing probability assessments, the device is effectively ‘betting against itself’, in the sense that each channel holds a different prior probability assessment over the signal space. This clearly could not happen if the device was not subjective, as a direct extension of Aumann’s famous common knowledge result (see [Aumann, 1976]).

Within this framework, we can formulate a device that is analogous to the generalized correlated equilibria in BDG.

Definition 6. A *generalized extended communication device (GECD)* is an ECD where for each i \mathcal{F}_i is a class of S such that:¹⁶

$$\bigcup_{i \in N} \mathcal{F}_i = S$$

¹⁴It is also in canonical form, but this is irrelevant for the purposes of this example.

¹⁵Clearly, the SECE constructed tends to a canonical CE as $\epsilon \rightarrow 0$.

¹⁶A class of a set X is a collection of subsets of X .

Clearly, this is an analogue of non-partitional information; the function

$$\begin{aligned} f_i : S &\rightarrow \mathcal{F}_i \\ f_i(s) &= f_{i,s} \end{aligned}$$

is the equivalent of a possibility correspondence as defined previously, whereas in the ECD case, f_i corresponds to an information function.

We are now ready to prove one of our central results, namely that every SECE has a decision-theoretically equivalent GECE counterpart. In order to this, we must first formally define the notion of decision-theoretic equivalence in this framework.

Definition 7. *Let*

$$d_e = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (\mathcal{F}_i)_{i \in N})$$

and

$$\hat{d}_e = ((\hat{S}_i)_{i \in N}, (\hat{M}_i)_{i \in N}, (\hat{\pi}_i)_{i \in N}, (\hat{\mathcal{F}}_i)_{i \in N})$$

be two extended communication devices. Then the two extended communication equilibria $(d_e, (\sigma, \delta))$ and $(\hat{d}_e, (\hat{\sigma}, \hat{\delta}))$ are **decision-theoretically equivalent** if there exist isomorphisms $\phi_i : \mathcal{F}_i \rightarrow \hat{\mathcal{F}}_i$ for each i , such that

1. $M = \hat{M}$
2. $\phi_i(f_{i,\sigma(t)}) = \hat{f}_{i,\hat{\sigma}(t)}$
3. $\pi_i(\mathcal{F}_i) = \hat{\pi}_i(\phi_i(\mathcal{F}_i))$
4. For $\hat{f}_{i,\cdot} = \phi_i(f_{i,\cdot})$,

$$\begin{aligned} &\sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \sum_{m \in M} \sum_{a \in A} \pi_i(m|f_{i,(r_i, \sigma_i(t_i))}) u_i(\nu_i(a_i|m_i, t_i), (\delta_{-i}(a_{-i}|m_{-i}, t_{-i}), \sigma(t), t)) \\ &= \\ &\sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \sum_{m \in M} \sum_{a \in A} \hat{\pi}_i(m|\hat{f}_{i,(\hat{r}_i, \hat{\sigma}_i(t_i))}) u_i(\nu_i(b_i|m_i, t_i), (\delta_{-i}(a_{-i}|m_{-i}, t_{-i}), \hat{\sigma}(t), t)) \end{aligned}$$

$$\forall r_i \in S_i, \forall \hat{r}_i \text{ s.t. } (\hat{r}_i, \hat{\sigma}_{-i}(t_{-i})) \in \phi_i((r_i, \sigma_{-i}(t_{-i}))) \forall b_i \in A_i, \forall \nu : M_i \times T_i \rightarrow \Delta A_i, \forall t_i \in T_i$$

In short, this states that strategies align on isomorphic information sets, and that expected utilities are equal. The isomorphism on information sets implies that the equilibria are behaviorally linked in the sense that, faced with the two information structures, the players/device make the same decisions. The first condition is necessary because players have fully certain on their respective outputs, i.e. on M_i . Note that the last condition is stronger than payoff equivalence, since it also enforces that isomorphic deviations from the equilibria yield equal expected utility. As noted in BDG, it is easily seen that decision-theoretic equivalence is an equivalence relation:

Remark 4. *Decision-theoretic equivalence is an equivalence relation on the set of SGECE.*

(*sketch of proof*). Transitivity follows from the composition of isomorphisms being an isomorphism. Symmetry follows from the inverse of an isomorphism being an isomorphism. Reflexivity is clear. \square

We can now present the central results of this section. In each case, after outlining the statement and proof, we discuss the intuition behind the proofs.

Theorem 1. *Let $d_e = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (\mathcal{F}_i)_{i \in N})$ be an SGECD, and $(d_e, (\sigma, \delta))$ the associated equilibrium of a game Γ_B . Then there exists a SECD $\hat{d}_e = ((\hat{S}_i)_{i \in N}, (\hat{M}_i)_{i \in N}, (\hat{\pi}_i)_{i \in N}, (\hat{\mathcal{F}}_i)_{i \in N})$ and strategies $(\hat{\sigma}, \hat{\delta})$ such that the associated SECE of Γ_B is decision-theoretically equivalent to $(d_e, (\sigma, \delta))$.*

Proof. Let $\hat{S}_i = \mathcal{F}_{\tau(i)}$, where $\tau : N \rightarrow N$ is bijective and has no fixed point.¹⁷ Then $\hat{S} = \mathcal{F}_{\tau(1)} \times \dots \times \mathcal{F}_{\tau(N)}$.

Set

$$\hat{\sigma}(t) = (f_{\tau(1), \sigma(t)}, \dots, f_{\tau(N), \sigma(t)})$$

For ease of notation, let $\hat{S} = \mathcal{F}_1 \times \dots \times \mathcal{F}_N$. Construct partitions on \hat{S} as follows: for $(R_1, \dots, R_N) \in \hat{S}$, let

$$\hat{\mathcal{P}}_i(R_1, \dots, R_N) = \{R_i\} \times \mathcal{F}_{-i}$$

Essentially, channel i is unable to distinguish between elements in \hat{S} that have the same i -th component. This induces the isomorphisms:

¹⁷For instance, let τ be the permutation $(1\ 2\ \dots\ N)$ on the group with N elements.

$$\begin{aligned}\phi_i : \mathcal{F}_i &\rightarrow \hat{\mathcal{F}}_i \\ \phi_i(f_i) &= \{f_i\} \times \mathcal{F}_{-i}\end{aligned}$$

Let the partitions $\hat{\mathcal{F}}_i$ be generated by ϕ_i , i.e. let

$$\begin{aligned}\hat{f}_{i,\hat{s}} &= \phi_i(\mathcal{P}_i(s)) \\ \hat{\mathcal{F}}_i &= \bigcup_{\hat{s} \in \hat{\mathcal{S}}} \hat{f}_{i,\hat{s}}\end{aligned}$$

Now construct the distributions:

$$\bar{\pi}_i(m|\hat{f}_{i,\hat{s}}) = \pi_i(m|s), \quad \text{where } \hat{f}_{i,\hat{s}} = \phi_i(\mathcal{P}_i(s))$$

Then $\hat{d}_e = ((\hat{S}_i)_{i \in N}, (\hat{M}_i)_{i \in N}, (\hat{\pi}_i)_{i \in N}, (\hat{\mathcal{F}}_i)_{i \in N})$ is the necessary device, and $(\hat{d}_e, (\hat{\alpha}, \delta))$ is decision-theoretically equivalent to $(d_e, (\sigma, \delta))$.

To see that this is indeed an equilibrium, it is clear that there is no incentive for any player to deviate from δ with respect to \hat{d}_B if there isn't with respect to d_B . To show that there is no incentive to deviate from $\hat{\sigma}(t)$, suppose that $\tau(j) = i$ and $\tau(i) = k$. Then

$$\begin{aligned}\hat{\mathcal{F}}_i(\hat{\sigma}(t)) &= \hat{\mathcal{F}}_i(\dots, f_{\tau(i),\sigma(t)}, \dots, f_{\tau(j),\sigma(t)} \dots) \\ &= \hat{\mathcal{F}}_i(\dots, \hat{s}_i, \dots, f_{i,\sigma(t)} \dots) \\ &= \hat{f}_{i,(\hat{s}_i, \hat{\sigma}_{-i}(t_{-i}))} \quad \forall \hat{s}_i \in \hat{S}_i,\end{aligned}$$

i.e. player i reports in the set $\mathcal{F}_{\tau(i)}$, hence any deviation by i will not change the i th component of $\hat{\sigma}(t)$, since τ has no fixed point, and will thus be in the same information set with respect to $\hat{\mathcal{F}}_i$. Hence, deviation at the input stage has no effect on each player's expected utility.

□

We prove the converse to this result.

Theorem 2. *Let $d_e = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (\mathcal{F}_i)_{i \in N})$ be an SECD, and $(d_e, (\sigma, \delta))$ the associated equilibrium of a game Γ_B . Then there exists a GECD $\hat{d}_e = ((\hat{S}_i)_{i \in N}, (\hat{M}_i)_{i \in N}, (\hat{\pi}_i)_{i \in N}, (\hat{\mathcal{F}}_i)_{i \in N})$ and*

strategies $(\hat{\sigma}, \hat{\delta})$ such that the associated GBCE of Γ_B is decision-theoretically equivalent to $(d_e, (\sigma, \delta))$.

Proof. Let $\hat{S} = S \times \{1, \dots, N\}$. Let $\hat{S}_{\underline{i}} = S_i \times \{1, \dots, N\}$ for some fixed $\underline{i} \in N$, and let $\hat{S}_i = S_i$ for $i \neq \underline{i}$. Construct $\hat{\sigma}(t)$ as

$$\begin{aligned}\hat{\sigma}_{-\underline{i}}(t) &= \sigma_{-\underline{i}}(t) \\ \hat{\sigma}_{\underline{i}}(t) &= (\sigma_{\underline{i}}(t), j), \quad \text{for any } j \in N\end{aligned}$$

Construct partitions on \hat{S} as follows: for $(s, j) \in \hat{S}$, let

$$\hat{\mathcal{P}}_i(s, j) = \mathcal{P}_i(s)$$

Then the isomorphisms ϕ_i naturally follow:

$$\begin{aligned}\phi_i : \mathcal{P}_i &\rightarrow \hat{\mathcal{P}}_i \\ \phi_i(R_i) &= R_i \times \{1, \dots, N\}\end{aligned}$$

Again, let the classes \mathcal{F}_i be generated by ϕ_i . Then construct the prior distribution:

$$\hat{\pi}(m|(s, i)) = \pi_i(m|s)$$

This distribution is well-defined because of the construction of the classes above, specifically, because the classes are pairwise disjoint, i.e.:

$$\mathcal{F}_i \cap \mathcal{F}_j = \emptyset, \quad \forall i, j \in N$$

Then $\hat{d}_e = ((\hat{S}_i)_{i \in N}, (\hat{M}_i)_{i \in N}, (\hat{\pi}_i)_{i \in N}, (\hat{\mathcal{F}}_i)_{i \in N})$ is the necessary GECD, and $(\hat{d}_e, (\hat{\alpha}, \delta))$ is the decision-theoretically equivalent equilibrium.

Again, it is clear that there is no incentive for any player to deviate from δ . To show that there is no incentive to deviate from $\hat{\sigma}(t)$, it suffices to note that in the construction of \hat{S} , the second component is essentially a dummy, in that it has no effect on either the partitions $\hat{\mathcal{F}}_i$ or the prior $\hat{\pi}$. Hence, player \underline{i} can report any element of the set $\{1, \dots, N\}$ without effect on the payoffs. Since the other strategies

are identical to the initial equilibrium, the no-deviation property is inherited.

□

In both the proofs above, the construction of the new signal space \hat{S} is the key step. Both constructions are broadly taken from the proofs of Propositions 4.1 and 4.2 in BDG. In Theorem 1, the construction is essentially expanding the signal space to allow each channel to be able to ‘imagine what all others know’, i.e. form the product space over all the classes of S , which represent the knowledge for each channel. This may seem like cheating; by effectively endowing each channel with the *combined* knowledge of all channels, it is easy to induce a partition on this construction, as has been shown. Indeed, the construction offers an implicit argument against the existence of non-partitional information; by extending the state-space sufficiently, one can ‘reveal’ a partition for a decision-maker, by allowing him to perceive what he can perceive. A similar observation is made by [Luce, 1956] in the context of preference theory.¹⁸ In Theorem 2, the construction is again natural in some sense. Essentially, by adding the dummy set $\{1, \dots, N\}$ to S , we are able to combine the given heterogenous priors into a common prior by indexing using this dummy set.

These observations are not intended to undermine the results. They simply serve to give the reader some insight into the choice of constructions, which, at first glance, may seem both arbitrary and mathematically recondite. Indeed, the objective of these results is simply to create a mapping between the two sets of equilibria.

3.1 Potential Drawbacks

Having introduced the concept of an ECD, we offer some potential shortcomings. Recall that the original aim of constructing such a device was to simulate existing communication devices whilst ‘opening up’ the process through which they generate output. As such, it should preferably retain as many properties of communication devices as possible, whilst explicating the ‘thinking’ process in as efficient a manner as possible. There are two ways in which extended communication devices fail by this measure.

First, the definition of the transition probability distributions is not strictly maintained by extended

¹⁸We thank Peter Hammond for pointing us toward this source.

devices; in a communication device, we have

$$Q_i : S \longrightarrow \Delta M,$$

whereas in an extended device,

$$\pi_i : \mathcal{F}_i \longrightarrow \Delta M,$$

Hence, the uncertainty embodied in the thinking process of an ECD is captured in the \mathcal{F}_i , rather than a direct thought process.

Second, and more importantly, extended devices do not perform Bayesian updating. In a sense, Bayesian updating is the most natural form of information processing available, operating through the following simple algorithm: prior understanding, observation, interpretation and subsequent understanding. Furthermore, it is the canonical model of information processing adopted within the game theory literature. Hence, it would be desirable to model the device's thinking process using Bayesian updating.

4 Bayesian Communication Equilibria

In the last section, we introduced the notion of an ECD, and discussed various possible drawbacks of such a device. We now attempt to address those drawbacks by constructing a new device, namely a *Bayesian communication device*. This new class of device both preserves the structure of the transition probabilities and incorporates Bayesian updating. We will see how extended devices represent a 'reduced-form' version of Bayesian devices, in the sense that, whilst it does not restrict the equilibrium payoff set, the device's thought process is far more simplistic.

Definition 8. A *Bayesian communication device (BCD)* is a tuple

$d_B = ((S_i)_{i \in N}, (M_i)_{i \in N}, \pi, q, (\mathcal{P}_i)_{i \in N})$, where S_i, M_i are as before, π is a prior distribution on S , q is a transition probability distribution from S to M , i.e. $q : S \rightarrow \Delta M$, and \mathcal{P}_i is an information function on S .

A *subjective Bayesian communication device (SBCD)* is a Bayesian communication device in which the prior π and the transition distribution q is allowed to vary for each i , i.e. a tuple

$d_B = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (q_i)_{i \in N}, (\mathcal{P}_i)_{i \in N})$.

A *generalized Bayesian communication device (GBCD)* is a Bayesian communication device in which \mathcal{P}_i is a possibility correspondence for each i .

A *subjective generalized Bayesian communication equilibrium (SGBCE)* of Γ_B is a Nash equilibrium of the game formed by attaching a subjective generalized Bayesian communication device to Γ_B .

Remark 5. Take a Bayesian game $\Gamma_B = (N, (T_i)_{i \in N}, p, (A_i)_{i \in N}, (u_i)_{i \in N})$, and a SGBCD $d_B = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (q_i)_{i \in N}, (\mathcal{P}_i)_{i \in N})$. Given a strategy profile (σ, δ) , the expected utility to player i of type t_i when playing the extended game Γ_B^d is

$$U_i((\sigma, \delta), t_i) = \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \sum_{m \in M} \sum_{s \in S} \sum_{a \in A} \bar{\pi}_i(s|\mathcal{P}_i(\sigma(t))) q_i(m|s) u_i((\delta(a|m, t), \sigma(t), t))$$

where $\bar{\pi}_i$ is channel i 's posterior, given by

$$\bar{\pi}_i(s'|\mathcal{P}_i(s)) = \begin{cases} \frac{\pi_i(s')}{\sum_{\hat{s} \in \mathcal{P}_i(s)} \pi_i(\hat{s})} & \text{if } s' \in \mathcal{P}_i(s) \\ 0 & \text{if } s' \notin \mathcal{P}_i(s) \end{cases}$$

Alternatively, let $Q_i : S \rightarrow \Delta M$ be defined as

$$Q_i(m|s) = \sum_{s' \in S} \bar{\pi}_i(s'|\mathcal{P}_i(s)) q_i(m|s')$$

Then

$$U_i((\sigma, \delta), t_i) = \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \sum_{m \in M} \sum_{a \in A} Q_i(m|\sigma(t)) u_i((\delta(a|m, t), \sigma(t), t))$$

The device works as follows. Having received an input vector, it performs Bayesian updating on the priors given the information encoded in the \mathcal{P}_i 's. It then combines this posterior with the transition distributions to generate an output. This is clearly a more natural extension of the communication device, as it preserves the existence of transition distributions from S to M ; in Remark 5, these is Q_i (we will call these distributions the *overall transition distributions*). The only modification to the standard communication device is in prescribing the manner in which the device explicitly interprets its inputs. Diagrammatically, we have:

$$\text{Communication device: } T \xrightarrow{\sigma} S \xrightarrow{\pi} M \xrightarrow{\delta} A$$

Bayesian communication device: $T \xrightarrow{\sigma} S \xrightarrow[\text{priors}]{\text{update}} S \xrightarrow{q} M \xrightarrow{\delta} A$

At first glance, it may seem as though, rather than dealing with the level of arbitrariness in the transition distributions, Bayesian devices simply add a layer of thinking, whilst maintaining the same arbitrariness in the distribution q . This is not the case. The process of Bayesian updating and the manner in which these posteriors are combined with the transition distribution results in a probabilistic weighting process; more weight is given to the distribution $q(M|s)$ corresponding to a greater posterior on s .

Remark 6. *If $\mathcal{P}_i(s) = s$ for all $s \in S$, and if $\pi_i = \pi$ and $q_i = q$ for all i , then the subjective Bayesian communication device $d_c = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (\mathcal{F}_i)_{i \in N})$ is a communication device.*

Moreover, we can show how extended devices are essentially Bayesian devices that operate using a simplified thought process.

Proposition 3. *For a given SGECD $d_e = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\hat{\pi}_i)_{i \in N}, (\mathcal{F}_i)_{i \in N})$, there exists a SGBCD $d_B = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (q_i)_{i \in N}, (\mathcal{P}_i)_{i \in N})$ such that:*

$$\hat{\pi}_i(m|f_{i,s}) = Q_i(m|s'), \quad \forall s' \in f_{i,s}, \quad \forall f_{i,s} \in \mathcal{F}_i, \quad \forall i \in N.$$

and vice versa.

Proof. The converse direction is clear; set $\mathcal{F}_i = \mathcal{P}_i$, and

$$\hat{\pi}_i(m|f_{i,s}) = Q_i(m|s'), \quad \forall s' \in f_{i,s}, \quad \forall f_{i,s} \in \mathcal{F}_i, \quad \forall i \in N.$$

as above.

Now Suppose $M = \{m_1, m_2, \dots, m_j\}$ and $S = \{s_1, s_2, \dots, s_k\}$. Then let the \mathbf{q}_i be the $k \times j$ right Markov transition matrix for d_B , i.e.

$$\mathbf{q}_i^T = \begin{pmatrix} q_i(m_1|s_1) & q_i(m_1|s_2) & \cdots & q_i(m_1|s_k) \\ q_i(m_2|s_1) & q_i(m_2|s_2) & \cdots & q_i(m_2|s_k) \\ \vdots & \vdots & \ddots & \vdots \\ q_i(m_j|s_1) & q_i(m_j|s_2) & \cdots & q_i(m_j|s_k) \end{pmatrix}$$

Let $\bar{\pi}_i(s)$ be the k -vector of channel i 's posterior on S , given $\mathcal{P}_i(s)$, and let

$$\bar{\Pi}_i = \begin{pmatrix} \bar{\pi}_i(s_1) & \bar{\pi}_i(s_2) & \cdots & \bar{\pi}_i(s_k) \end{pmatrix}$$

Then, if

$$\mathbf{Q}_i^T = \begin{pmatrix} Q_i(m_1|s_1) & Q_i(m_1|s_2) & \cdots & Q_i(m_1|s_k) \\ Q_i(m_2|s_1) & Q_i(m_2|s_2) & \cdots & Q_i(m_2|s_k) \\ \vdots & \vdots & \ddots & \vdots \\ Q_i(m_j|s_1) & Q_i(m_j|s_2) & \cdots & Q_i(m_j|s_k) \end{pmatrix}$$

we have by Remark 5 that

$$\mathbf{Q}_i^T = \mathbf{q}_i^T \bar{\Pi}_i$$

Now let $q_i(m|s) = \pi_i(m|f_{i,s})$, let $\mathcal{P}_i(s) = f_{i,s}$ and let the prior π be uniform over S . Then

$$(\bar{\Pi}_{i'})_{ij} = \begin{cases} \frac{1}{|\mathcal{P}_{i'}(s)|} & \text{if } s_i \in \mathcal{P}_{i'}(s_j) \\ 0 & \text{otherwise} \end{cases}$$

Thus,

$$\begin{aligned} \left(\mathbf{q}_{i'}^T \bar{\Pi}_{i'} \right)_{ij} &= \sum_{s_k \in \mathcal{P}_{i'}(s_j)} \bar{\pi}_{i'}(s_k | \mathcal{P}_{i'}(s_j)) q_{i'}(m_i | s_k) \\ &= \sum_{s_k \in \mathcal{P}_{i'}(s_j)} \frac{q_{i'}(m_i | s_k)}{|\mathcal{P}_{i'}(s_j)|} \\ &= \pi_{i'}(m_i | f_{i,s}), \quad \text{where } s_j \in f_{i,s} \end{aligned}$$

□

This shows that we can think of extended devices as ‘reduced-form’ examples of Bayesian devices. We will see that such an analogy has heightened resonance in the context of the proofs of Theorems 3 and 4, which are essentially enriched versions of those of Theorems 1 and 2.

Like subjective/generalized extended communication devices, subjective/generalized Bayesian commu-

nication devices can generate equilibrium payoffs that pareto-dominate communication devices:¹⁹

Example. Let $\Gamma_B = (\{1, 2\}, (T_i), (p_i), (A_i), (u_i))$ be a Bayesian game: $T_1 = \{t_1, s_1\}$, $T_2 = \{t_2, s_2\}$, $p_1(s_1) = p_1(t_1) = \frac{1}{2}$, $p_2(s_2) = p_2(t_2) = \frac{1}{2}$, $A_1 = \{T, B\}$, $A_2 = \{L, R\}$, and von-Neumann-Morgenstern payoffs are given in the tables below:

$(t_1, t_2), (s_1, s_2)$	L	R
T	0,8	3,3
B	1,1	0,0

$(t_1, s_2), (s_1, t_2)$	L	R
T	0,0	3,3
B	1,1	8,0

i.e., the first payoff structure obtains in states $(t_1, t_2), (s_1, s_2)$, and the second obtains in states $(t_1, s_2), (s_1, t_2)$. We now define a device d_B for Γ_B . Let $S = \{(t_1, t_2), (s_1, t_2), (t_1, s_2), (s_1, s_2)\} \equiv \{1, 2, 3, 4\}$.

Let $M = \{m_1, m_2, m_3, m_4\}$. Define the $(\mathcal{P}_i)_{i \in I}$ by the following:

$$\begin{aligned}
 \mathcal{P}_1(1) &= \{1, 2\} & \mathcal{P}_2(1) &= \{1, 3, 4\} \\
 \mathcal{P}_1(2) &= \{1, 2, 3\} & \mathcal{P}_2(2) &= \{2, 3, 4\} \\
 \mathcal{P}_1(3) &= \{1, 2, 3\} & \mathcal{P}_2(3) &= \{2, 3, 4\} \\
 \mathcal{P}_1(4) &= \{1, 2, 4\} & \mathcal{P}_2(4) &= \{4, 3\}
 \end{aligned}$$

Let $\pi(s) = \frac{1}{4}$ for all $s \in S$. Let $q(m_i|S) = \mathbf{1}_i$, where $\mathbf{1}_i$ is a vector of 0's and a 1 in the i -th component.

Now define the strategy profiles as:

$$\begin{aligned}
 \alpha_i(t_i) &= t_i \quad \forall i \in N, t_i \in T_i \\
 \delta_1(m_1, t_1) &= U, \quad \forall m_1 \in M_1, t_1 \in T_1 \\
 \delta_2(m_2, t_2) &= R, \quad \forall m_2 \in M_2, t_2 \in T_2
 \end{aligned}$$

Then $(d_B, (\alpha, \delta))$ is a GBCE of Γ_B . For further discussion, see Appendix A.

¹⁹This example is based on Example 2 in BDG.

We now redefine decision-theoretic equivalence in this setting:

Definition 9. *Let*

$$d_B = ((S_i)_{i \in N}, (M_i)_{i \in N}, \pi, q, (\mathcal{P}_i)_{i \in N})$$

and

$$\hat{d}_B = \left((\hat{S}_i)_{i \in N}, (\hat{M}_i)_{i \in N}, \hat{\pi}, \hat{q}, (\hat{\mathcal{P}}_i)_{i \in N} \right)$$

be two SGBCDs. Then the two SGBCE's $(d_c, (\sigma, \delta))$ and $(\hat{d}_c, (\hat{\sigma}, \hat{\delta}))$ are **decision-theoretically equivalent** if their exist isomorphisms $\phi_i : \mathcal{P}_i \rightarrow \hat{\mathcal{P}}_i$ for each i , such that

1. $M = \hat{M}$
2. $\phi_i(\mathcal{P}_i(\sigma(t))) = \hat{\mathcal{P}}_i(\hat{\sigma}(t))$
3. For $\hat{\mathcal{P}}_i(\cdot) = \phi_i(\mathcal{P}_i(\cdot))$,

$$\begin{aligned} & \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \sum_{m \in M} \sum_{s \in S} \sum_{a \in A} \bar{\pi}_i(s|\mathcal{P}_i(r_i, \sigma_{-i}(t_{-i}))) q_i(m|s) u_i \left(\nu_i(a_i|m_i, t_i), (\delta_{-i}(a_{-i}|m_{-i}, t_{-i}), \sigma(t), t) \right) \\ & = \\ & \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \sum_{m \in M} \sum_{\hat{s} \in \hat{S}} \sum_{a \in A} \bar{\hat{\pi}}_i(\hat{s}|\hat{\mathcal{P}}_i(\hat{r}_i, \hat{\sigma}_{-i}(t_{-i}))) \hat{q}_i(m|\hat{s}) u_i \left(\nu_i(b_i|m_i, t_i), (\delta_{-i}(a_{-i}|m_{-i}, t_{-i}), \hat{\sigma}(t), t) \right) \\ & \quad \forall r_i \in S_i, \forall \hat{r}_i \text{ s.t. } (\hat{r}_i, \hat{\sigma}_{-i}(t_{-i})) \in \phi_i((r_i, \sigma_{-i}(t_{-i}))) \quad \forall b_i \in A_i, \forall \nu : M_i \times T_i \rightarrow \Delta A_i, \forall t_i \in T_i \end{aligned}$$

We now prove the central result of this section. Again, the proof borrows elements from the proof of Proposition 4.1 in BDG, but also makes some amendments.

Theorem 3. *Let $d_B = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (q_i)_{i \in N}, (\mathcal{P}_i)_{i \in N})$ be an SGBCD, and $(d_B, (\sigma, \delta))$ the associated equilibrium of a game Γ_B . Then there exists a SBCD $\hat{d}_B = \left((\hat{S}_i)_{i \in N}, (\hat{M}_i)_{i \in N}, (\hat{\pi}_i)_{i \in N}, (\hat{q}_i)_{i \in N}, (\hat{\mathcal{P}}_i)_{i \in N} \right)$ and strategies $(\hat{\sigma}, \hat{\delta})$ such that the associated SBCE of Γ_B is decision-theoretically equivalent to $(d_B, (\sigma, \delta))$.*

Proof. Let $\mathcal{S}_i = \mathcal{P}_i(S)$, i.e. the partition of S generated by \mathcal{P}_i . Let $\hat{S}_i = \mathcal{S}_{\tau(i)}$, where $\tau : N \rightarrow N$ is bijective and has no fixed point. Then $\hat{S} = \mathcal{S}_{\tau(1)} \times \dots \times \mathcal{S}_{\tau(N)}$.

Set

$$\hat{\sigma}(t) = (\mathcal{P}_{\tau(1)}(\sigma(t)), \dots, \mathcal{P}_{\tau(N)}(\sigma(t)))$$

Again, for ease of notation, let $\hat{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_N$. Construct the partitions on \hat{S} as follows: for $(R_1, \dots, R_N) \in \hat{S}$, let

$$\hat{\mathcal{P}}_i(R_1, \dots, R_N) = \{R_i\} \times \mathcal{S}_{-i}$$

Essentially, channel i is unable to distinguish between elements in \hat{S} that have the same i -th component. This induces the set of isomorphisms:

$$\begin{aligned} \phi_i : \mathcal{P}_i &\rightarrow \hat{\mathcal{P}}_i \\ \phi_i(R_i) &= \{R_i\} \times \mathcal{S}_{-i} \end{aligned}$$

Now construct the posteriors:

$$\bar{\pi}_i(R_1, \dots, R_N | \{\bar{R}_i\} \times \mathcal{S}_{-i}) = \pi_i(\{s \in S : \mathcal{P}_j(s) = R_j, i \neq j | \bar{R}_i\})$$

The left-hand side of this expression is the probability of the element (R_1, \dots, R_N) occurring, given that the element $(\dots, \bar{R}_i, \dots)$ has occurred. It is essentially a projection/counting process; identify the signals that lie in the intersection of the information sets R_1, \dots, R_N , for some fixed \bar{R}_i and then calculate the probability of this event occurring.

The input strategy profile $\hat{\sigma}(t)$ and the partitions $\hat{\mathcal{P}}_i$ satisfy part 2 of the definition for decision-theoretic equivalence. It remains to construct the transition distributions $\hat{q}_i : \hat{S} \rightarrow M$. For payoff equivalence to hold, we require that

Condition 1. $\hat{Q}_i(m|\hat{s}) = Q_i(m|s)$, where $\hat{\mathcal{P}}_i(\hat{s}) \in \phi_i(\mathcal{P}_i(s))$

In words, we require that the overall transition distributions are equal on isomorphic information sets.

Let

$$\hat{q}_i(m|\hat{s}) = \frac{1}{|\hat{\mathcal{P}}_i(\hat{s})|} \left(\frac{Q_i(m|s)}{\bar{\pi}_i(\hat{s}|\hat{\mathcal{P}}_i(\hat{s}))} \right), \quad \forall m \in M, \hat{s} \in \hat{S}, \text{ where } \hat{\mathcal{P}}_i(\hat{s}) \in \phi(\mathcal{P}_i(s))$$

Then

$$\begin{aligned}
\hat{Q}_i(m|\hat{s}) &= \sum_{s \in \hat{S}} \bar{\pi}_i(s|\hat{\mathcal{P}}_i(\hat{s})) \hat{q}_i(m|s) \\
&= \sum_{s \in \hat{\mathcal{P}}_i(\hat{s})} \bar{\pi}_i(s|\hat{\mathcal{P}}_i(\hat{s})) \hat{q}_i(m|s) \\
&= \sum_{s \in \hat{\mathcal{P}}_i(\hat{s})} \frac{\bar{\pi}_i(s|\hat{\mathcal{P}}_i(\hat{s})) Q_i(m|s)}{\bar{\pi}_i(s|\hat{\mathcal{P}}_i(s)) |\hat{\mathcal{P}}_i(\hat{s})|} \\
&= Q_i(m|s)
\end{aligned}$$

The second equality of this derivation holds by definition of the posterior, the third holds by direct substitution for $\hat{q}_i(m|\hat{s})$ and the fourth because $\bar{\pi}_i(s|\hat{\mathcal{P}}_i(s)) = \bar{\pi}_i(s|\hat{\mathcal{P}}_i(\hat{s}))$ for all $s \in \mathcal{P}_i(\hat{s})$.²⁰

Thus, the device $\hat{d}_B = \left((\hat{S}_i)_{i \in N}, (M_i)_{i \in N}, (\hat{\pi}_i)_{i \in N}, (\hat{q}_i)_{i \in N}, (\hat{\mathcal{P}}_i)_{i \in N} \right)$ is the necessary device, and the equilibrium $(\hat{d}_B, (\hat{\sigma}, \hat{\delta}))$ is the necessary equilibrium.

That this profile is an equilibrium follows an identical argument to that in the proof of Theorem 1. \square

We now prove the converse to the Theorem 3.

Theorem 4. *Let $d_B = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (q_i)_{i \in N}, (\mathcal{P}_i)_{i \in N})$ be an SBCD, and $(d_B, (\sigma, \delta))$ the associated equilibrium of a game Γ_B . Then there exists a GBCD $\hat{d}_B = \left((\hat{S}_i)_{i \in N}, (\hat{M}_i)_{i \in N}, \hat{\pi}, \hat{q}, (\hat{\mathcal{P}}_i)_{i \in N} \right)$ and strategies $(\hat{\sigma}, \hat{\delta})$ such that the associated GBCE of Γ_B is decision-theoretically equivalent to $(d_B, (\sigma, \delta))$.*

Proof. Let $\hat{S} = S \times \{1, \dots, N\}$. Let $\hat{S}_{\underline{i}} = S_i \times \{1, \dots, N\}$ for some fixed $\underline{i} \in N$, and let $\hat{S}_i = S_i$ for $i \neq \underline{i}$. Construct $\hat{\sigma}(t)$ as

$$\begin{aligned}
\hat{\sigma}_{-\underline{i}}(t) &= \sigma_{-\underline{i}}(t) \\
\hat{\sigma}_{\underline{i}}(t) &= (\sigma_{\underline{i}}(t), j), \quad \text{for any } j \in N
\end{aligned}$$

Construct the partitions on \hat{S} as follows: for $(s, j) \in \hat{S}$, let

$$\hat{\mathcal{P}}_i(s, j) = (\mathcal{P}_i(s), i)$$

²⁰It is important to note that this last step only holds because the $\hat{\mathcal{P}}_i$'s are partitions.

Then the isomorphisms ϕ_i naturally follow:

$$\begin{aligned}\phi_i : \mathcal{P}_i &\rightarrow \hat{\mathcal{P}}_i \\ \phi_i(R_i) &= (\{R_i\}, i)\end{aligned}$$

Now construct the prior distribution:

$$\bar{\pi}(s, j) = \frac{1}{G} \pi_i(s | \mathcal{P}_i(s)),$$

where $G = \sum_{(s,j) \in \hat{S}} \pi_i(s | \mathcal{P}_i(s))$ is the normalization constant. Let $\hat{q}(m | (s, i)) = q_i(m | s)$ for all i . Then Condition 1 is satisfied, and the strategy profile $(\hat{\sigma}, \delta)$ is an equilibrium of Γ_B with respect to \hat{d}_B .

Again, that this profile is in equilibrium follows from the proof of Theorem 2. □

For numerical examples of the constructions used in Theorems 1 and 2, see Appendices B and C respectively. It is important to note that the GBCE constructed in the last proof does not satisfy non-delusion.²¹ Figure 1 demonstrates this observation graphically.

²¹This result mirrors that following Proposition 4.2 in BDG, since the signal-space construction above largely mirrors their construction of $\tilde{\Omega}$.




Figure 1: Chain of Inclusions generated by Theorems 3 and 4

NND refers to ‘non-non-delusion’, i.e. information structures that do not satisfy non-delusion.

Remark 7. *Non-delusion does not expand the set of GBCEs of a Bayesian game Γ_B .*

Proof. Follows directly from Remark 4 and sequential application of Theorems 3 and 4. □

Owing to the vast array of different devices, equilibria, associated payoff spaces and relationships between these objects, it is worth summarizing the results obtained graphically. See Figure 2.

5 Further Topics

5.1 The Revelation Principle and Generalized Bayesian Devices

In section 2.3, we outlined the notions of canonical devices, equilibria and the Revelation Principle. The system of linear inequalities defining a canonical equilibrium were shown to not only embody a computational tractable problem, but also to provide a characterization of the space of transition distributions that qualify as canonical, i.e. $\mu : T \rightarrow \Delta A$ is an honest-and-obedient canonical equilibrium of Γ_B if and only if

$$\sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \sum_{a \in A} \pi(a|t) u^i(t, a) \geq \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \sum_{a \in A} \pi(a|s_i, t_{-i}) u^i(t, f(a_i), a_{-i}),$$

$$\forall i \in N, \forall t_i, s_i \in T_i, \forall f : A_i \rightarrow A_i$$

Essentially, the Revelation Principle comprises of two parts. First, it defines and classifies the set of honest-and-obedient canonical equilibria. Second, it maps non-canonical equilibria into this set. However, we have shown through examples how the set of equilibrium payoffs is expanded by allowing for subjective priors. Since a canonical equilibrium generated through the Revelation Principle clearly

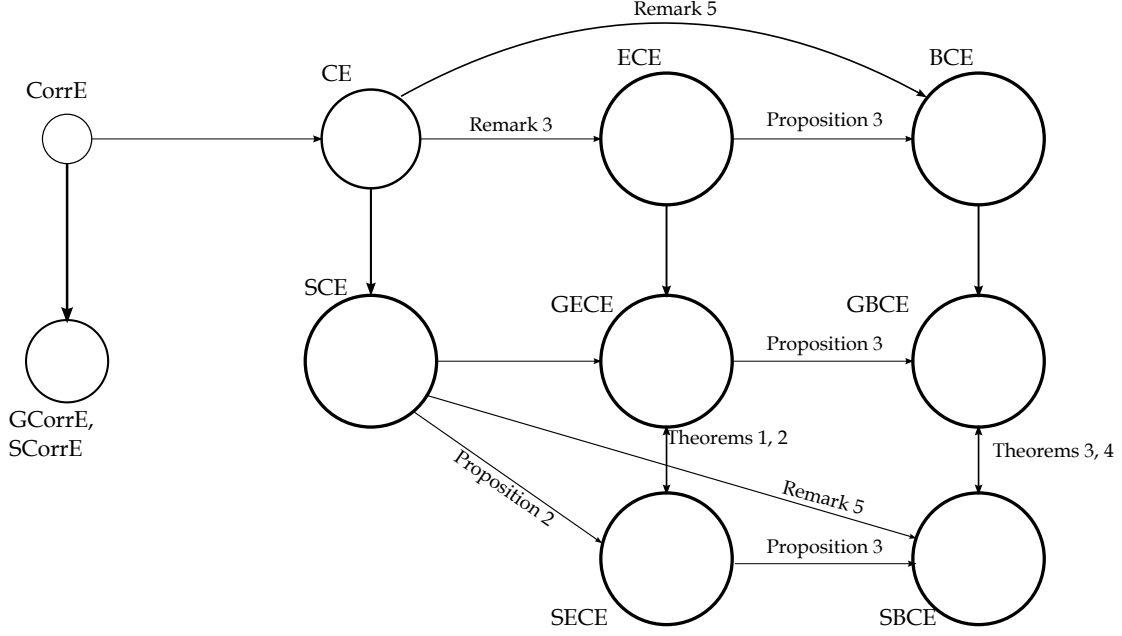


Figure 2: Graphical Summary of Results

cannot contain payoff profiles outside the set of non-canonical equilibria payoff profiles, it is evident that, if we are to assert the Revelation Principle in the context of either subjective or generalized devices, we must re-define the set of honest-and-obedient canonical equilibria, i.e. we must relax the constraints embodied in the linear equations above for a canonical device μ to qualify as a generalized or subjective canonical device. The question is thus embodied in the following diagram:

$$\begin{array}{ccc}
 CE's & \xrightarrow{RevPrinc} & CCE's \\
 \downarrow & & \downarrow \\
 SCE's, GCE's & \xrightarrow{RevPrinc} & ?
 \end{array}$$

First, let us consider subjective devices, be they extended, Bayesian or orthodox communication devices. Just as was the case for correlated equilibria, it is easy to prove the Revelation Principle in this setting.

Theorem 5. Let $(d_e, (\sigma, \delta))$ be an SECE with respect to the device

$d_e = ((S_i)_{i \in N}, (M_i)_{i \in N}, (\pi_i)_{i \in N}, (\mathcal{F}_i)_{i \in N})$ for the game $\Gamma_B = (N, (T_i)_{i \in N}, p, (A_i)_{i \in N}, (u_i)_{i \in N})$. Then there is a canonical subjective communication equilibrium $(d_c^c, (\sigma_c, \delta_c))$ for Γ_B that is payoff equivalent to $(d_e, (\sigma, \delta))$, in which every player truthfully reveals his type, and takes the action prescribed to them by the device.

Proof. If such a canonical equilibrium exists, it is by construction the case that σ_c^i and δ_c^i are the identity maps on T_i and M_i respectively, for each i . It remains to construct a family of transition distributions:

$$\mu_i : T \longrightarrow A$$

such that $(d_c^e, (\sigma_c, \delta_c))$ is a canonical communication equilibrium for Γ_B that is payoff-equivalent to $(d_e, (\sigma, \delta))$, i.e.

$$\sum_{t \in T} p(t_{-i}|t_i) \sum_{m \in M} \pi_i(m|f_{i,\sigma(t)}) u_i((\sigma, \delta), t) = \sum_{t \in T} p(t_{-i}|t_i) \sum_{a \in A} \mu_i(a|t) u_i((t, a), t)$$

Choose

$$\mu_i(a|t) = \sum_{(\sigma, \delta)} \sum_{m \in \delta^{-1}(a)} \left(\prod_{i \in N} \lambda_i(\sigma_i, \delta_i|t_i) \right) \pi_i(m|f_{i,\sigma(t)})$$

where the distribution λ_i is the naturally induced marginal for each i and t_i , and $\delta^{-1}(a) = \{m \in M : \delta(a|m) > 0\}$.

The device constructed works as follows:

1. The device asks players to report their types.
2. The device simulates the reports (the $\sigma_i(t_i)$ s) that would have been sent according to the original equilibrium, given the types reported in step 1.
3. The device determines both the distribution on M that would have been arrived at under the original device, for each i , and the action profile players would have chosen under the original device given this output.
4. The device prescribes this action.

It is now clear why the device-strategy profile combination created is an equilibrium; if there was any incentive to deviate under the new equilibrium, there must have been in the original equilibrium, which is absurd. □

The statement and proof of the Revelation Principle for subjective Bayesian devices is identical in nature, with the expressions for payoff equivalence duly altered.

Let us now consider generalized devices. To tackle this question, we draw analogy to Proposition 5.1 in BDG. This result characterizes the set of honest-and-obedient canonical generalized correlated

equilibria of a normal-form game Γ , and hence verifies the first part of the Revelation Principle.²² We now propose a classification of all honest-and-obedient canonical generalized Bayesian communication equilibria.

Definition 10. *Define*

$$Q_\mu(a_i, t_i) = \left\{ q : T_{-i} \rightarrow \Delta A_{-i} \mid \text{supp } q(s_{-i}) \subset \text{supp } \mu(\cdot | a_i, (s_{-i}, t_i)) \forall s_{-i} \in T_{-i}, \right. \\ \left. \sum_{t \in T} p(t_{-i} | t_i) \sum_{a_{-i} \in A_{-i}} q(a_{-i} | t_{-i}) u_i((t, a), t) \geq \sum_{t \in T} p(t_{-i} | t_i) \sum_{a_{-i} \in A_{-i}} q(a_{-i} | t_{-i}) u_i((t, b_i, a_{-i})) \right\}$$

Then $Q_\mu(a_i, t_i)$ is the set of all transition distributions on A_{-i} with support contained in $\mu(\cdot | a_i, t_i)$ for which a_i and t_i are optimal for i . For an arbitrary set Y , let

$$\text{aff } Y = \left\{ \sum_n \beta_n y_n \mid y_n \in Y, \sum_n \beta_n = 1 \right\}$$

Conjecture 1. $\mu : T \rightarrow \Delta A$ is a canonical GBCE if and only if $\mu(\cdot | a_i, t_i) \in \text{aff } Q_\mu(a_i, t_i)$

Remark 8. The statement that $\mu(\cdot | a_i, t_i) \in Q_\mu(a_i, t_i)$ is precisely that $\mu(\cdot | a_i, t_i)$ is a canonical BCE.

The form of Conjecture 1 mirrors that of Proposition 5.1 in BDG, with the definition of $Q_\mu(a_i, t_i)$ extended to cover transition distributions, rather than simply probability distributions. Remark 8 shows that the conditions for a canonical GBCE are broader than for a canonical BCE, which is logical. Proving it is beyond the scope of this project, as it would require extensive construction and discussion in and of itself.

5.2 Fallible Talk and Bayesian Communication Equilibria

In Section 2.4, we discussed how direct communication between players can be formalized into a game-theoretic solution concept, namely a *cheap-talk extension* of Γ_B , and how the equilibrium payoffs generated by such systems are contained in those that use private communication, examples of which include the various notions of communication equilibria employed throughout this paper. We can now construct a form of cheap talk that embodies the ‘public’ counterpart to BCDs. Consider the following rules of an extension to Γ_B :

²²They refer to these equilibria as generalized correlated equilibrium distributions.

1. Players learn their types.
2. Given these types, players send a message, or signal, to some subset of the remaining players.
3. Each player can interpret these messages up to some level of uncertainty, represented by a prior and an information function on the signal space.
4. Having interpreted these signals, the original game is played.

It is clear how this scheme is closely related to that followed in a BCE.²³ Such a scheme clearly has practical application. One can think of a situation in which a person solicits the advice of an expert, but cannot fully comprehend their advice, maybe for linguistic or technical reasons. Such communication could be termed *fallible talk*. Note that this applies whether the players have partitional or non-partitional information. More specifically,

Definition 11. *Let Γ_B be a Bayesian game. We define a **fallible talk extension** of Γ_B as a game in which many interim phases of unmediated communication are allowed before Γ_B is played. More precisely, let S_i be a set of signals for player i , $\pi_i(\cdot|s_i)$ be a prior measure on S for i , and \mathcal{P}_i be an information function on S for i . At each stage, every player selects an $s_i \in S_i$. These choices are then revealed to a subset of the remaining players. Having received the vector of inputs s_{-i} , player i combines this with s_i and interprets this new information by updating $\pi_i(\cdot|s_i)$ using \mathcal{P}_i in the following fashion: suppose player i receives messages from players in $J \subset N$. Then he updates his prior conditional on the event E_i , where*

$$E_i = \{(s_J, s_{-J}), \forall s_{-J} \in S_{-J}\}$$

At the end of the communication phase, the game is played.

Of course, such communication makes more sense in the canonical case, i.e. where each the signals each player can send lie in their type space.

It is important to note that [Green, Stokey, 1980] also allow for subjective priors and non-partitional information in their cheap talk games. However, they assume these properties for the underlying game, rather than the additional communication phases, i.e. they allow the prior p on T to be different for each player, and allow the agent to learn his type in a non-partitional fashion. With fallible talk, the

²³One difference is that there is no message space involved in the scheme above, but this difference is purely cosmetic. We could easily insert such an interim phase, but this would add no descriptive power to the specification of the process.

underlying game involves full rationality and a common prior, whilst the communication phase may not.

An interesting application of fallible talk is to the *opinion game*.²⁴ To describe this game, we must abandon the type space construction of a Bayesian game, and use a state-space/partition construction, as discussed in Section 2.1.

Definition 12. *The Opinion Game* Let $\Omega \subset \mathbb{R}$, let $p(\Omega)$ be a prior distribution, let \mathcal{P}_i be a finite information function,²⁵ let $A_i = \mathbb{R}$, and let $u_i(a, \omega) = -(a_i - \omega)^2$ (the players are rewarded for approximating the true state ω with greater accuracy). Then $\Gamma_O = (\{1, 2\}, \Omega, \pi, \mathcal{P}_i, A_i, u_i)$ is a Bayesian game.

Now, consider a fallible talk extension of Γ_O in which the signal space for each player is Ω , and the priors and partitions are the same as for the underlying game. Then we have:

Remark 9. For the fallible talk extension of the opinion game described above, there exists an equilibrium in which each player takes the same action.

This is simply a re-wording of Aumann’s famous common knowledge result; if two people share a common prior, and their posteriors are common knowledge, then their posteriors are equal. Specifically, it follows the form outlined in [Geanakoplos, Polemarchakis, 1982]; rather than common knowledge, they assume a finite number of revisions of the players’ posteriors through direct communication. The result holds even if we use possibility correspondences that satisfy non-delusion and balancedness. This is essentially Theorem 6 in [Geanakoplos, 1989].

6 Conclusion

The purpose of this paper was to extend the notion of a communication equilibrium to allow for explicit modeling of the underlying device’s thinking process. We achieved this through the introduction of two new devices, namely *extended* and *Bayesian* communication devices. We then showed how, in both cases, allowing the device to adopt ‘multiple priors’ is equivalent to imposing bounded rationality on its information processing capabilities. We discussed the Revelation Principle in the context of extended and Bayesian equilibria, and concluded by discussing an analogue of cheap talk, namely *fallible talk* in the context of Bayesian communication equilibria.

²⁴As outlined in [Geanakoplos, 1989].

²⁵i.e. $|\mathcal{P}_i(\Omega)| < \infty$.

The remit of this paper restricts the breadth and depth to which we can explore the concepts introduced in it. As such, there remain many topics worth exploring.

For instance, it may be beneficial to consider refinements to the solution concepts involved. Nash equilibrium is effectively the only solution concept used throughout; for all types of device, we look at the Nash equilibria of the games formed by attaching the device. Indeed, parts of the proofs of Theorems 1,2,3 and 4 would clearly not hold for more refined solution concepts.²⁶ If we interpret the game as being static, then ex-post and ex-ante equilibria agree, so Nash equilibrium is sufficient as a solution concept. If, however, we interpret the extended game as being in extensive form, this solution concept clearly runs into problems, as exemplified by such refinements as sub-game perfection and sequential equilibrium à la introduced in [Kreps, Wilson, 1982]. [Myerson, 1986] suggests an analogue to sequential equilibrium for communication equilibria, namely *sequential communication equilibrium*. The concepts discussed informally in Sections 5.1 and 5.2 could be formalized. It would be ideal to prove Conjecture 1, and discuss its implications at length. Regarding fallible talk, it would be beneficial to not only formalize the concept and results pertaining to it, but also to apply the concept to existing examples, e.g. from the principal-agent literature. One could, for example, take the classic quadratic-uniform game presented in [Crawford, Sobel, 1982], and apply fallible talk, so that principal now receives the agent's report, but has difficulty understanding it. This may well alter the structure of the set of so-called 'partition equilibria' associated with the game.

We have been primarily interested in characterizing the set of equilibria generated by a class of communication devices, rather than fixing a device and examining the equilibria associated with it. It would be interesting to consider the effects on welfare of marginal improvements in the device's information (the fineness of the partitions/classes), for any fixed SECE, GECE, SBCE, GBCE, etc. This type of analysis was performed in [Green, Stokey, 1980], who find that, in their class of cheap talk games, improvements in information do not necessarily generate welfare improvements.

Finally, it would be of interest to find direct, meaningful applications of the concepts involved, in particular BCDs. This paper is largely theoretic, and the few examples contained in it have been purposefully simplistic. Possible applications include to stylized signaling games (see [Forges, 1990]) and market games à la [Shapley, Shubik, 1969].

²⁶For instance, the no-deviation reasoning used for $\hat{\sigma}$ in all cases.

References

- [Aumann, 1967] Aumann, R., 1967, *Subjectivity and Correlation in Randomized Strategies*, Journal of Mathematical Economics, vol 1, pp 67-96.
- [Aumann, 1976] Aumann, R., 1976, *Agreeing to Disagree*, The Annals of Statistics, vol 4, no 6, pp 1236-1239.
- [Aumann, 1987] Aumann, R., 1987, *Correlated Equilibrium as an Expression of Bayesian Rationality*, Econometrica, vol 55, no 1, pp 1-18.
- [Crawford, Sobel, 1982] Crawford, V., Sobel, J., 1982, *Strategic Information Transmission*, Econometrica, vol 50, no 6, pp 1431-1451.
- [Forges, 1985] Forges, F., 1985, *Correlated Equilibria in a Class of Repeated Games with Incomplete Information*, International Journal of Game Theory, vol 14, pp 129-150.
- [Forges, 1986] Forges, F., 1986, *An Approach to Communication Equilibria*, Econometrica, vol 54, pp 1375-1385.
- [Forges, 2009] Forges, F., 2009, *Correlated Equilibrium and Communication in Games*, in Myers, R.A., 2009, *Encyclopedia of Complexity and Systems Science*, Springer.
- [Forges, 1990] Forges, F., 1990, *Equilibria with Communication in a Job Market Example*, Quarterly Journal of Economics, vol 105, No 2, pp 375-398.
- [Geanakoplos, 1989] Geanakoplos, J., 1989, *Game Theory Without Partitions, and Applications to Speculation and Consensus*, Cowles Foundation Discussion Paper, no 914, pp 1375-1385.
- [Geanakoplos, Polemarchakis, 1982] Geanakoplos, J., Polemarchakis, H., 1982, *We Cannot Disagree Forever*, Journal of Economic Theory, vol 28, pp 192-200.
- [Green, Stokey, 1980] Green, J.R., Stokey, N.L., 1980, *A Two-Person Game of Information Transmission*, Discussion Papers 418, Northwestern University, Center for Mathematical Studies in Economics and Management Science.
- [Harsanyi, 1967] Harsanyi, J. C., 1967, *Games with Incomplete Information Played by Bayesian Players*, Parts I, II, III, Management Science, vol 14, pp 159-182, 320-334, 486-502.

- [Kreps, Wilson, 1982] Kreps, D.M., Wilson, R., 1982, *Sequential Equilibria*, *Econometrica*, vol 50, pp 863-894.
- [Lehrer, 1996] Lehrer, E., 1996, *Mediated Talk*, *International Journal of Game Theory*, vol 25, pp 177-188.
- [Luce, 1956] Luce, R.D., 1956, *Semioorders and a Theory of Utility Discrimination*, *Econometrica*, vol 24, no 2, pp 178-191.
- [Milgrom, Stokey, 1982] Milgrom, P., Stokey, N., 1982, *Information, Trade and Common Knowledge*, *Journal of Economic Theory*, vol 26, pp 17-27.
- [Myerson, 1982] Myerson, R., 1982, *Optimal Coordination Mechanisms in Generalized Principal-Agent Problems*, *Journal of Mathematical Economics*, vol 10, pp 67-81.
- [Myerson, 1986] Myerson, R., 1986, *Multistage Games with Communication*, *Econometrica*, vol 54, no 2, pp 323-358.
- [Morris, 1995] Morris, S., 1995, *The Common Prior Assumption in Economic Theory*, *Economics and Philosophy*, vol 11, pp 227-253.
- [Osborne, Rubinstein, 1994] Osborne, M., Rubinstein, A., 1994, *A Course in Game Theory*, MIT Press.
- [Samet, 1990] Samet, D., 1990, *Ignoring Ignorance and Agreeing to Disagree*, *Journal of Economic Theory*, vol 21, pp 192-207.
- [Shapley, Shubik, 1969] Shapley, L., Shubik, M., 1969, *On Market Games*, *Journal of Economic Theory*, vol 1, pp 9-25.
- [Shin, 1993] Shin, H.S., 1993, *Logical Structure of Common Knowledge*, *Journal of Economic Theory*, vol 60, pp 1-13.

A Appendix A

First, consider the example from Section 3. The honest-and-obedient strategy profile is in equilibrium if

$$\sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \sum_{a \in A} \pi(a|t) u^i(t, a) \geq \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \sum_{a \in A} \pi(a|r_i, t_{-i}) u^i(t, f(a_i), a_{-i}),$$

$$\forall i \in N, \forall t_i, r_i \in T_i, \forall f : A_i \rightarrow A_i$$

We verify the inequalities for player 1 of type s_1 . There are four cases, since $|T_1| = |A_1| = 2$. First, note that 1's expected utility in equilibrium is given by

$$\begin{aligned} U_1((s_1, a_1|s_1)) &= p(t_2|s_1) \left(\mu_1((T, L)|s_1) u_1(s_1, (T, L)) + 0 + \mu_1((B, R)|s_1) u_1(s_1, (B, R)) + 0 \right) \\ &= \frac{1}{2} \left(\frac{1+\epsilon}{2} \cdot 1 + \frac{1-\epsilon}{2} \cdot (-1) \right) \\ &= \frac{\epsilon}{2} \end{aligned}$$

1. $\mathbf{r_1 = s_1, f(T) = T}$

Then 1's expected utility becomes

$$U_1((s_1, f(a_1)|s_1)) = \frac{1}{2} \left(\frac{1+\epsilon}{2} \cdot 1 + \frac{1-\epsilon}{2} \cdot (-1) \right) = \frac{\epsilon}{2}$$

Similarly,

2. $\mathbf{r_1 = t_1, f(T) = T}$

$$U_1((t_1, f(a_1)|s_1)) = \frac{1}{2} \left(\frac{1+\epsilon}{2} \cdot (-1) + \frac{1-\epsilon}{2} \cdot 1 \right) = -\frac{\epsilon}{2}$$

3. $\mathbf{r_1 = s_1, f(T) = B}$

$$U_1((s_1, f(a_1)|s_1)) = \frac{1}{2} \left(\frac{1-\epsilon}{2} \cdot 1 + \frac{1+\epsilon}{2} \cdot (-1) \right) = -\frac{\epsilon}{2}$$

4. $\mathbf{r_1 = t_1, f(T) = B}$

$$U_1((t_1, f(a_1)|s_1)) = \frac{1}{2} \left(\frac{1-\epsilon}{2} \cdot (-1) + \frac{1+\epsilon}{2} \cdot 1 \right) = \frac{\epsilon}{2}$$

From these calculations, it is clear that the inequalities hold. Similar computations are possible for the remaining types and players.

Next, consider the example from Section 4. Each player receives an expected utility of 3 in equilibrium. Thus, there is no incentive for player 1 to deviate. The posterior matrices are given by

$$\bar{\Pi}_1 = \begin{pmatrix} \frac{1}{2} & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & 0 & 0 & \frac{1}{3} \end{pmatrix} \quad \bar{\Pi}_2 = \begin{pmatrix} \frac{1}{3} & 0 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & 0 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & \frac{1}{2} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & \frac{1}{2} \end{pmatrix}$$

Since $\mathbf{q}_i = \mathbf{I}_4$, we have that $\mathbf{Q}_i = \bar{\Pi}_i$. Now, consider player 1. By lying, he received a payoff of at most $\frac{8}{3}$ (by also deviating from T to B), which is less than 3. A symmetric argument holds for player 2.

B Appendix B

We construct numerical examples to illustrate the constructions used in Theorems 1. Throughout this appendix, we adopt the matrix notation used in Proposition 3. Since the action-profile δ is invariant under the transformation, we will omit this from the following discussion, i.e. we will assume payoffs and a δ exist such that the remainder of the construction yields an equilibrium.

Now suppose we have a Bayesian game with two players, where $T_1 = \{t_1, s_1\}$, $T_2 = \{t_2\}$, so $T = \{(t_1, t_2), (s_1, t_2)\} \equiv \{t, s\}$, $p_1(s_1) = p_1(t_1) = \frac{1}{2}$, $p_2(s_2) = p_2(t_2) = \frac{1}{2}$, $A_1 = A_1$, $A_2 = A_2$, and payoffs are simply $u_i(a, t)$.

Take the following device: $S = T \equiv \{1, 2\}$ (i.e $t = 1$ and $s = 2$), $M = \{A, B\}$, the priors given by

$$\begin{aligned} \pi_1(1) &= \frac{1}{2} & \pi_2(1) &= \frac{1}{3} \\ \pi_1(2) &= \frac{1}{2} & \pi_2(2) &= \frac{2}{3}, \end{aligned}$$

the possibility correspondences given by

$$\begin{aligned} \mathcal{P}_1(1) &= \{1, \} := a & \mathcal{P}_2(1) &= \{1, 2\} := c \\ \mathcal{P}_1(2) &= \{1, 2, \} := b & \mathcal{P}_2(2) &= \{2\} := d, \end{aligned}$$

and the transition matrices given by

$$\mathbf{q}_1^T = \begin{pmatrix} \frac{1}{2} & \frac{2}{3} \\ \frac{1}{2} & \frac{1}{3} \end{pmatrix} \quad \mathbf{q}_2^T = \begin{pmatrix} \frac{1}{3} & \frac{1}{2} \\ \frac{2}{3} & \frac{1}{2} \end{pmatrix}$$

Then the posterior matrices are given by

$$\bar{\mathbf{\Pi}}_1 = \begin{pmatrix} 1 & \frac{1}{2} \\ 0 & \frac{1}{2} \end{pmatrix} \quad \bar{\mathbf{\Pi}}_2 = \begin{pmatrix} \frac{1}{3} & 0 \\ \frac{2}{3} & 1 \end{pmatrix}$$

so the overall transition matrices are given by

$$\begin{aligned} \mathbf{Q}_1^T &= \begin{pmatrix} \frac{1}{2} & \frac{2}{3} \\ \frac{1}{2} & \frac{1}{3} \end{pmatrix} \begin{pmatrix} 1 & \frac{1}{2} \\ 0 & \frac{1}{2} \end{pmatrix} \\ &= \begin{pmatrix} \frac{1}{2} & \frac{7}{12} \\ \frac{1}{2} & \frac{5}{12} \end{pmatrix} \end{aligned}$$

and

$$\begin{aligned} \mathbf{Q}_2^T &= \begin{pmatrix} \frac{1}{3} & \frac{1}{2} \\ \frac{2}{3} & \frac{1}{2} \end{pmatrix} \begin{pmatrix} \frac{1}{3} & 0 \\ \frac{2}{3} & 1 \end{pmatrix} \\ &= \begin{pmatrix} \frac{4}{9} & \frac{1}{2} \\ \frac{5}{9} & \frac{1}{2} \end{pmatrix} \end{aligned}$$

Now take the report profile α to consist of the identity map for each i , and assume that the combination of this profile, with the above SGBCD, plus an arbitrary δ is in equilibrium. We proceed to construct the decision-theoretically equivalent SBCE via the construction used in the proof of Theorem 1. Hence, we have

$$\hat{S}_1 = \{c, d\}, \hat{S}_2 = \{a, b\},$$

Note that there is only one τ that satisfies the necessary conditions, i.e. $\tau(1) = 2$, $\tau(2) = 1$. We construct $\hat{\sigma}$ by

$$\begin{aligned}\hat{\sigma}(t) &= (\mathcal{P}_2(\sigma(t)), \mathcal{P}_1(\sigma(t))) \\ &= (\mathcal{P}_2(1), \mathcal{P}_1(1)) \\ &= (c, a)\end{aligned}$$

Similarly, $\hat{\sigma}(s) = (d, b)$. Again, for convenience, write

$$\hat{S} = \{a, b\} \times \{c, d\} = \{(a, c), (a, d), (b, c), (b, d)\}$$

Then the partitions $\hat{\mathcal{P}}_i$ are given by

$$\begin{aligned}\hat{\mathcal{P}}_1(a, c) = \hat{\mathcal{P}}_1(a, d) &= \{(a, c), (a, d)\} & \hat{\mathcal{P}}_1(b, c) = \hat{\mathcal{P}}_1(b, d) &= \{(b, c), (b, d)\} \\ \hat{\mathcal{P}}_2(a, c) = \hat{\mathcal{P}}_2(b, c) &= \{(a, c), (b, c)\} & \hat{\mathcal{P}}_2(a, d) = \hat{\mathcal{P}}_2(b, d) &= \{(a, d), (b, d)\},\end{aligned}$$

The posterior matrices are given by

$$\tilde{\mathbf{\Pi}}_1 = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}, \quad \tilde{\mathbf{\Pi}}_2 = \begin{pmatrix} \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ \frac{2}{3} & 0 & \frac{2}{3} & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}$$

For example, suppose (a, c) was reported. Then

$$\begin{aligned}\hat{\pi}_1\left((a, c) \mid \{(a, c), (a, d)\}\right) &= \pi_1\left(s \in S, \mathcal{P}_2(s) = \{1, 2\} \mid \{1\}\right) \\ &= 1 \\ \hat{\pi}_1\left((a, d) \mid \{(a, c), (a, d)\}\right) &= \pi_1\left(s \in S, \mathcal{P}_2(s) = \{2\} \mid \{1\}\right) \\ &= 0 \\ \hat{\pi}_1\left((b, c) \mid \{(a, c), (a, d)\}\right) &= 0 \\ \hat{\pi}_1\left((b, d) \mid \{(a, c), (a, d)\}\right) &= 0\end{aligned}$$

Next, we calculate the $\hat{\mathbf{q}}_i$'s as being

$$\hat{\mathbf{q}}_1^T = \begin{pmatrix} \frac{1}{2} & \frac{2}{3} & \frac{1}{2} & \frac{2}{3} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{2} & \frac{1}{3} \end{pmatrix} \quad \hat{\mathbf{q}}_2^T = \begin{pmatrix} \frac{1}{3} & \frac{1}{2} & \frac{1}{3} & \frac{1}{2} \\ \frac{2}{3} & \frac{1}{2} & \frac{2}{3} & \frac{1}{2} \end{pmatrix}$$

Thus, the overall transition matrices are given by

$$\begin{aligned} \hat{\mathbf{Q}}_1^T &= \begin{pmatrix} \frac{1}{2} & \frac{2}{3} & \frac{1}{2} & \frac{2}{3} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{2} & \frac{1}{3} \end{pmatrix} \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix} \\ &= \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & \frac{7}{12} & \frac{7}{12} \\ \frac{1}{2} & \frac{1}{2} & \frac{5}{12} & \frac{5}{12} \end{pmatrix} \end{aligned}$$

and

$$\begin{aligned} \hat{\mathbf{Q}}_2^T &= \begin{pmatrix} \frac{1}{3} & \frac{1}{2} & \frac{1}{3} & \frac{1}{2} \\ \frac{2}{3} & \frac{1}{2} & \frac{2}{3} & \frac{1}{2} \end{pmatrix} \begin{pmatrix} \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ \frac{2}{3} & 0 & \frac{2}{3} & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix} \\ &= \begin{pmatrix} \frac{4}{9} & \frac{1}{2} & \frac{4}{9} & \frac{1}{2} \\ \frac{5}{9} & \frac{1}{2} & \frac{5}{9} & \frac{1}{2} \end{pmatrix} \end{aligned}$$

By inspection, Condition 1 holds.

C Appendix C

We now construct an example for Theorem 2. Take the same underlying game Γ_B as in the previous. For the SBCD, take S , M , π_i , q_i and σ also as previously, but for convenience, relabel $S = \{\alpha, \beta\}$. Let the partitions be given by

$$\begin{aligned}\mathcal{P}_1(\alpha) &= \{\alpha, \beta\} & \mathcal{P}_2(\alpha) &= \{\alpha\} \\ \mathcal{P}_1(\beta) &= \{\alpha, \beta\} & \mathcal{P}_2(\beta) &= \{\beta\},\end{aligned}$$

Combining the transition distributions with the posteriors yields

$$\begin{aligned}\mathbf{Q}_1^T &= \begin{pmatrix} \frac{1}{2} & \frac{2}{3} \\ \frac{1}{2} & \frac{1}{3} \end{pmatrix} \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \\ &= \begin{pmatrix} \frac{7}{12} & \frac{7}{12} \\ \frac{5}{12} & \frac{5}{12} \end{pmatrix}\end{aligned}$$

and

$$\begin{aligned}\mathbf{Q}_2^T &= \begin{pmatrix} \frac{1}{3} & \frac{1}{2} \\ \frac{2}{3} & \frac{1}{2} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \\ &= \begin{pmatrix} \frac{1}{3} & \frac{1}{2} \\ \frac{2}{3} & \frac{1}{2} \end{pmatrix}\end{aligned}$$

We now construct the associated GBCD. Then

$$\hat{S} = \{\alpha, \beta\} \times \{1, 2\} = \{(\alpha, 1), (\alpha, 2), (\beta, 1), (\beta, 2)\}$$

Suppose $\hat{i} = 2$, then $\hat{S}_1 = \{t_1, s_1\}$, $\hat{S}_2 = \{t_2\} \times \{1, 2\}$. Define $\hat{\sigma}$ by

$$\begin{aligned}\hat{\sigma}_1(t) &= \sigma_1(t) = t_1 \\ \hat{\sigma}_2(t) &= (\sigma_2(t), 2), \text{ say}\end{aligned}$$

Similarly for s . The $\hat{\mathcal{P}}_i$'s are given by:

$$\begin{aligned}\hat{\mathcal{P}}_1(\alpha, 1) &= \hat{\mathcal{P}}_1(\alpha, 2) = \hat{\mathcal{P}}_1(\beta, 1) = \hat{\mathcal{P}}_1(\beta, 2) = \{(\alpha, 1), (\beta, 1)\} \\ \hat{\mathcal{P}}_2(\alpha, 1) &= \hat{\mathcal{P}}_2(\alpha, 2) = \{(\alpha, 2)\} & \hat{\mathcal{P}}_2(\beta, 1) &= \hat{\mathcal{P}}_2(\beta, 2) = \{(\beta, 2)\},\end{aligned}$$

The prior distribution $\hat{\pi}$ is calculated by

$$\begin{aligned}\hat{\pi}(\alpha, 1) &= \frac{1}{G} \pi_1(\alpha | \mathcal{P}_1(\alpha)) = \frac{1}{3} \times \frac{1}{2} = \frac{1}{6} \\ \hat{\pi}(\alpha, 2) &= \frac{1}{3} \\ \hat{\pi}(\beta, 1) &= \frac{1}{6} \\ \hat{\pi}(\beta, 2) &= \frac{1}{3}\end{aligned}$$

Hence, the posterior matrices become:

$$\tilde{\mathbf{\Pi}}_1 = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & 0 \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \tilde{\mathbf{\Pi}}_2 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix}$$

The transition matrix $\hat{\mathbf{q}}$ is given by

$$\hat{\mathbf{q}}^T = \begin{pmatrix} \frac{1}{2} & \frac{1}{3} & \frac{2}{3} & \frac{1}{2} \\ \frac{1}{2} & \frac{2}{3} & \frac{1}{3} & \frac{1}{2} \end{pmatrix}$$

Hence the overall transitions become

$$\begin{aligned}\hat{\mathbf{Q}}_1^T &= \begin{pmatrix} \frac{1}{2} & \frac{1}{3} & \frac{2}{3} & \frac{1}{2} \\ \frac{1}{2} & \frac{2}{3} & \frac{1}{3} & \frac{1}{2} \end{pmatrix} \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & 0 \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & 0 \end{pmatrix} \\ &= \begin{pmatrix} \frac{7}{12} & \frac{7}{12} & \frac{7}{12} & \frac{7}{12} \\ \frac{5}{12} & \frac{5}{12} & \frac{5}{12} & \frac{5}{12} \end{pmatrix}\end{aligned}$$

and

$$\begin{aligned}\hat{\mathbf{Q}}_2^T &= \begin{pmatrix} \frac{1}{2} & \frac{1}{3} & \frac{2}{3} & \frac{1}{2} \\ \frac{1}{2} & \frac{2}{3} & \frac{1}{3} & \frac{1}{2} \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix} \\ &= \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{2} & \frac{1}{2} \\ \frac{2}{3} & \frac{2}{3} & \frac{1}{2} & \frac{1}{2} \end{pmatrix}\end{aligned}$$

Condition 1 clearly holds by inspection.