The scope of provability*

Esteban Peralta† and Fernando Tohmé‡

November 8, 2020

Abstract

We explore the relationship between evidence and knowledge, when knowledge is described by a partition of a finite state space and evidence is represented by a collection of sets of messages that is measurable with respect to the partition. We show that there is evidence for an event only if the event is self-evident—i.e., known at every state. Thus, an event is public—commonly known at all of its states—whenever there is *mutual* evidence for it. It follows that the notion of *mutual evidence* captures the infinite regress embedded in the notion of common knowledge. We show that these results offer a novel foundation for understanding knowledge and common knowledge, implying that total provability is possible only if the partition is either the coarsest or the finest one. We also show that the existence of partial provability outside these two extremes hinges on the non-monotonic nature of provability.

Keywords: Knowledge - Evidence - Common Knowledge - Monotonicity.

1 Introduction

When the set of messages that an agent can send varies with her information, the agent can offer *evidence* of, or *prove*, some of what she knows. This context, often referred as a situation of partial provability [14], contrasts with settings in which messages convey no information and only cheap-talk communication is possible [4]. Since an agent's ability to provide evidence limits her capacity to misrepresent her information, the study of models with partial provability has attracted a considerable amount of attention in the analyses of communication and mechanism design problems.[1]

Despite the importance of partial provability, its relationship with knowledge has not been fully analyzed. This note uncovers the relationship between an agent's evidence and knowledge in a model in which knowledge is described by a partition of a finite state space and evidence is represented by a collection of sets of messages that is measurable with respect to the partition [13]. Thus, we do not assume that agents are necessarily informed about the actual state but that they know every event they can prove.

We proceed by providing a simple, yet powerful result (Proposition 1). We show that an agent can only prove *self-evident* events—i.e., events that are known at each one of its states (see, e.g., [3]). Our main goal is to state and discuss the implications of this result.[2]

Proposition 1 provides a sufficient condition for common knowledge since it implies that an event is public [16] among a group of agents—commonly known by each of them at each of its states—whenever the event is *mutually* proved by each agent. This consequence has, in turn, two important epistemic takeaways. On the one hand, it implies that mutual evidence is enough to capture the infinite regress embedded in the notion of common knowledge. On the other, it suggests a relevant connection with the well known relationship between common knowledge and coordination. To see this, suppose that two agents privately observe both a signal and a set of messages that varies with the signals they receive. Each agent is deciding whether to invest on a project that is only successful if the economic conditions are "healthy", and so each decides to join if and only if it is a commonly knowledge between them that the economic conditions are indeed healthy.

*We are grateful to Navin Kartik, Shaowei Ke, Frederic Koessler and Larry Samuelson for very helpful discussions. The usual disclaimer applies.

†Department of Economics and Stephen M. Ross School of Business, University of Michigan.

‡Departamento de Economía, Universidad Nacional del Sur and CONICET.

[1] For studies in communication contexts, see., e.g., [5], [9], [15] and [6]. For mechanism design problems, see, e.g., [8], [1], [12], [2] and [11].

[2] This point is also made in [13], but its consequences are not studied.

What we show is that this common knowledge arises whenever *both have evidence* to support the fact that the economic conditions are healthy. Thus, the presence of evidence might solve coordination problems without communication.

The second consequence of Proposition 1 is a novel foundation for the notions of knowledge and common knowledge. Since the algebra generated by a partition embeds every evidence structure that is measurable with respect to the partition [13], the algebra can be interpreted as describing the class of events that can be proved, *in principle*. That is, the events that are *provable* in the sense that there exists evidence for them in *some* evidence structure measurable with respect to the partition. Thus, Proposition 1 implies that an event is self-evident *if and only if* it is provable. This unveils a new characterization of knowledge and common knowledge. Namely, an event is known if and only if it is implied by a provable one. Hence, an event is public if and only if it is *mutually provable* and commonly known if and only if it is implied by a mutually provable event. We view these characterizations as suggesting that knowledge and common knowledge can be seen, not as primitives, but as *emerging* from evidence.

Finally, Proposition 1 sheds light on the extent of an agent's proof ability. Given that the literature has spanned the whole range of proof abilities, from cheap talk communication to total provability—situations in which an agent can prove every event that she knows (see, e.g., [17], [7] and [10])—we believe that uncovering the scope of evidence might offer a better understanding of its implications. We show that Proposition 1 entails that total provability is equivalent to the requirement that the agent's ability to prove is monotonic. We show, however, that this is only possible when the partition is either the coarsest or the finest one. Thus, an agent is able to prove every event she knows only a the top and bottom of the lattice of partitions, or equivalently if either she is always informed about the state she is at, or she cannot distinguish any state from another.

A natural follow-up question is when *partial* provability is possible if the agent is not informed about the state but holds nontrivial information. We introduce a joint condition on evidence and knowledge, termed *size-measurability*, that requires that, for any two states, the number of events an agent knows increases (weakly) with the number of events that she can prove. We then show that, when an agent's partition is not the finest one, she can prove *some* event she knows if and only if either her proof ability is not monotonic or size-measurability fails. Finally, we show that, when the class of every state in an agent's partition is a non-singleton, partial provability is possible if and only if her proof ability is not monotonic. These results highlight that monotonicity of evidence plays a dual role; namely, it is necessary for total provability but, within a large class of environments, also necessary for cheap-talk.

The plan of the paper is as follows. Section 2 introduces the notions of knowledge and evidence, as well as the fundamental condition that evidence must vary with respect to knowledge. Section 3 introduces the main claim of the paper, and Section 4 discusses its main consequences. Finally, Section 5 characterizes the existence of partial provability.

2 The environment

2.1 Information

Let S be a finite set of states of the world and $I = \{1, ..., n\}$ a finite set of n agents. We denote by $i \in I$ a generic agent and by Π_i is i's *information structure*, i.e. a partition of S. $\Pi_i(s)$ is the collection of states that agent i considers possible when the state is s. An information structure for I is a collection $\Pi := \{\Pi_i\}_{i \in I}$, where Π_i is the information structure of agent i.

Fix an information structure Π. Agent i knows event $E \subseteq S$ at s if $\Pi_i(s) \subseteq E$. Thus, the set of states at which agent i knows an event E is defined as:

$$K_i(E) := \{s \in S : \Pi_i(s) \subseteq E\}.$$

For any $E \subseteq S$, let $K_I(E)$ denote the event that "everyone knows" E; i.e.:

$$K_I(E) := \bigcap_{i \in I} K_i(E).$$

Recursively define the order $m \geq 1$ "everyone knows" event, K_I^m, by $K_I^m(E) = K_I(K_I^{m-1}(E))$, where $K_I^1(E) = K_I(E)$. Thus, define the event E is common knowledge, $CK(E)$, as follows:

$$CK(E) := \bigcap_{m=1}^{\infty} K_I^m(E).$$

Common knowledge can be equivalently described in terms of the existence of public events.[3] To describe this equivalence, more terminology is needed.

An event E is self-evident for i if i knows E at every one of its states; i.e., if $E \subseteq K_i(E)$. On the other hand, since by definition $K_i(E) \subseteq E$ for every i and every E, we can equivalently state that an event E is self-evident for i if $K_i(E) = E$. Notice that any event that is self-evident for i can be written as the union of elements of i's partition. That is, for any agent i the set of events that are self-evident for i, $\mathcal{A}_{\Pi_i}$, is given by the algebra generated by Π_i, minus the empty set. That is,

$$\mathcal{A}_{\Pi_i} := \{E \subseteq S : E = \bigcup_{s' \in E} \Pi_i(s')\}.$$

For any s, $\mathcal{A}_{\Pi_i}(s) := \{E \in \mathcal{A}_{\Pi_i} : s \in E\}$. An event E is **public** at s if it is evident knowledge for every i at s; i.e., if $E \in \mathcal{A}_{\Pi_i}(s)$ for every i.[4]

2.2 Evidence

Let $\mathcal{E}_i$ denote a finite set of messages for agent i, and let $e_i \in \mathcal{E}_i$ denote a generic message. We denote by $M_i(s) \subseteq \mathcal{E}_i$ the set of messages that i has available when the state is s and define $M_i =: \bigcup_{s \in S} M_i(s)$. Let $\wp(X)$ denote the power-set of a set X. The map $\mathcal{I}_i : \mathcal{E}_i \to \wp(S)$ gives the "interpretation" – or informational content – of every message possessed by i. Indeed, when i sends message $e_i \in M_i$ she *proves*, or offers *evidence*, that the true state must lie in the event:

$$\mathcal{I}_i(e_i) = \{s \in S : e_i \in M_i(s)\}.$$

Hence, we can identify each message e_i with the event $\mathcal{I}_i(e_i)$. We will then let $\mathbf{M}_i(s)$ denote the set of *events* that agent i can prove when the state is s; i.e.:

$$\mathbf{M}_i(s) := \{E \subseteq S : E = \mathcal{I}_i(e_i) \text{ for some } e_i \in M_i(s)\}.$$

We say that $\mathbf{M}_i := \bigcup_{s \in S} \mathbf{M}_i(s)$ is an **evidence structure** for agent i. Define $\mathbf{M} := \{\mathbf{M}_i\}_{i \in I}$. Given $\mathbf{M}$, we say that there is **mutual evidence** for an event E at state s if $E \in \mathbf{M}_i(s)$ for every $i \in I$.

By construction, evidence structures satisfy two properties. First, for any s and any $E \in \mathbf{M}_i(s)$, we have that $s \in E$. That is, no agent can prove an event that is false. Second, evidence is *consistent* in the following sense: $E \in \mathbf{M}_i(s)$ implies $E \in \mathbf{M}_i(s')$ for every $s' \in E$. Intuitively, consistency means that if an agent can prove event E in some state, she must be able to prove it at every state where the event is true. Indeed, proving an event requires a message that, by definition, must be available at every state where the event is true.

2.3 Models

A **model for agent** i is a pair $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ where $\mathbf{M}_i$ is measurable with respect to Π_i; i.e., for every s and s' such that $\Pi_i(s) = \Pi_i(s')$, we have that $\mathbf{M}_i(s) = \mathbf{M}_i(s')$. In every model $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ we have that $E \in \mathbf{M}_i(s) \Rightarrow \Pi_i(s) \subseteq E$ for every E.

Notice that, by construction, $\mathcal{A}_{\Pi_i}$ is an evidence structure that is measurable with respect to Π_i. Thus, we can write $\mathcal{M}_{\mathcal{A}_{\Pi_i}} := (\Pi_i, \mathcal{A}_{\Pi_i})$.

3 Evidence and self-evidence

The next proposition, that claims that an agent can only provide evidence of events that are self-evident for her, constitutes the main result of the paper:

Proposition 1. *For every $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ and every s, $\mathbf{M}_i(s) \subseteq \mathcal{A}_{\Pi_i}(s)$.*

[3] Alternatively, common knowledge can be described in terms of the meet of Π, i.e. the finest common coarsening of Π.

[4] Thus, public events can be written as the union of elements of the meet of Π.

Proof. Fix any agent i and suppose, contrary to hypothesis, that there exists a model $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$, a state s, and an event $E \in \mathbf{M}_i(s)$ such that $E \notin \mathcal{A}_{\Pi_i}(s)$. Then, there exists some $s' \in E$ such that $\Pi_i(s') \nsubseteq E$. Yet consistency entails that $E \in \mathbf{M}_i(s')$, contradicting that $\mathbf{M}_i$ is measurable with respect to Π_i. □

Intuitively, Proposition 1 shows that consistency strengthens the requirement that an agent's evidence must be measurable with respect to her knowledge. Indeed, Proposition 1 says that an agent is not only restricted to offer evidence of what she knows, but is in fact restricted to offer evidence of what it is self-evident to her.[5] The rest of the paper is devoted to state and discuss the implications of Proposition 1.

4 Main implications

4.1 Coordination and common knowledge

The first consequence of Proposition 1 involves the set of public events:

Corollary 1. *For every $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ and every s, if $E \in \bigcap_i \mathbf{M}_i(s)$, then $E \in \bigcap_i \mathcal{A}_{\Pi_i}(s)$; i.e., an event E is public at s whenever there is mutual evidence for E at s.*

Corollary 1 gives a sufficient condition for (evident) common knowledge. That is, a group of individuals commonly know an event if everyone of them has evidence for the event. In our view, the importance of this corollary lies on its connection with the well known relationship between coordination and common knowledge. Corollary 1 suggests that the existence of mutual evidence is all that might be needed to solve coordination problems. Corollary 1 also implies that the infinite regress embedded in the notion of common knowledge is captured by the simpler notion of mutual evidence.

4.2 A foundation for knowledge and common knowledge

For any agent i, the algebra generated by the partition Π_i, $\mathcal{A}_{\Pi_i}$, can be seen as the agent's "canonical" evidence structure in the sense that it constitutes the richest evidence structure that is measurable with respect to Π_i. It follows that, for a given partition, the set of events that the agent can prove, in principle—namely those that she proves in *some* measurable evidence structure—coincides with the set of her self-evident events. We formalize this idea with the following notion:

Definition 1. *An event E is **provable** given Π_i if there exists a model $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ and a state s such that $E \in \mathbf{M}_i(s)$.*

Notice that this definition fixes a partition but varies the evidence structure within the set of those that are measurable with respect to the partition. The following result is immediate:

Proposition 2. *An event E is provable given Π_i if and only if $E \in \mathcal{A}_{\Pi_i}$.*

Proof. The *if* part is obvious since $\mathcal{M}_{\mathcal{A}_{\Pi_i}}$ is a model. For the *only if* part, take any event E such that there exists some $\mathbf{M}_i$ and some state s with $E \in \mathbf{M}_i(s)$. Since $\mathbf{M}_i(s) \subseteq \mathcal{A}_{\Pi_i}(s)$ for every s by Proposition 1, it follows that $E \in \mathcal{A}_{\Pi_i}(s)$. □

This result strengthens Proposition 1 by showing that its converse holds when we focus our attention on the class of provable events. Given the well known characterizations of knowledge and common knowledge in terms of self-evident events (see, e.g., [3]), Proposition 2 entails that knowledge and common knowledge can be characterized, respectively, in terms of what is provable and mutually provable:

Corollary 2. *An event is known at s by i if and only if it is implied by an event that is provable by i at s.*

Corollary 3. *An event is public at s if and only if it is mutually provable at s. i.e., provable by every i. Thus, an event is commonly known at s if and only if it is implied by a mutually provable event at s.*

[5] Koessler (2004) also noticed that any evidence structure measurable with respect to a partition must be, state by state, a subset of the set of self-evident events. Yet he did not study the consequences of this property, which constitute the main goal of this paper.

The first statement in Corollary 3 implies that the class of events that are *not* public are precisely those that cannot be proved, *in principle*, by *some* agent. Conversely, every event that *can be* proved by all agents, although not necessarily in every measurable evidence structure, must be evident common knowledge among them.

The second statement in Corollary 3 provides a foundation for common knowledge in terms of the set of events for which there can be mutual evidence. We see this statement, together with Corollary 2, as suggesting that knowledge and common knowledge emerge from what is provable.

4.3 Total Provability

Proposition 2 implies that a known event is *not* provable if and only if it intersects an element of the partition that is not contained in it. This sheds light on the class of situations in which an agent's knowledge is provable in the sense that, for every E and s such that $\Pi_i(s) \subseteq E$, we have that $E \in \mathcal{A}_{\Pi_i}(s)$. These situations are often referred to as situations of *total provability* (see, e.g., [17], [7] and [10]). That is, total provability requires an agent's knowledge and evident knowledge to coincide.

In this subsection, we show that total provability is generally impossible. To see this, we need first to define some concepts:

Definition 2. *A model $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is one of* ***total provability*** *if, for every s and every $E \subseteq S$ such that $\Pi_i(s) \subseteq E$, we have that $E \in \mathbf{M}_i(s)$.*

Total provability models are models in which an agent's knowledge is evidence-based in the sense that is restricted by her proof-ability. Given that evidence is required to be measurable with respect to knowledge, total provability models are those in which the agent knows an event if and only if she is able to prove it. In total provability models a agent's proof ability is monotonic:

Definition 3. *$\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is* ***monotonic*** *(M) if, for every s:*

$$E \in \mathbf{M}_i(s) \Rightarrow F \in \mathbf{M}_i(s) \text{ for every } F \text{ such that } E \subseteq F.$$

At the same time, total provability implies that an agent is able to prove every signal she receives:

Definition 4. *$\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is* ***information-certifiable*** *(IC) if $\Pi_i(s) \in \mathbf{M}_i(s)$ for every s.*

This condition extends the notion of complete provability in [14] to situations in which the agent is not necessarily informed about the state, and is closely related to the notion of own-type certifiability proposed in [9]. Total provability can be characterized in terms of (M) and (IC), being equivalent to the claim that the algebra generated by the partition is closed under supersets:

Corollary 4. *Fix any agent i. The following statements are equivalent:*

1. *$\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is a model of total provability;*
2. *$\mathcal{M}_i$ satisfies (M) and (IC);*
3. *$\mathcal{A}_{\Pi_i}(s)$ is closed under supersets for every s.*

Proof. We first show that 1. and 2. are equivalent. Suppose that a model $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ satisfies both (IC) and (M). The former implies that $\Pi_i(s) \in \mathbf{M}_i(s)$ for every s. By (M), it follows that $F \in \mathbf{M}_i(s)$ for every s and every F such that $\Pi_i(s) \subseteq F$. Thus, $\mathcal{M}_i$ is a model of total provability. To show that 1. implies 2., notice that if $\mathcal{M}_i$ is a model of total provability, then $\Pi_i(s) \in \mathbf{M}_i(s)$ for every s. Hence, $\mathcal{M}_i$ is (IC). To show that $\mathcal{M}_i$ satisfies (M) it is sufficient to consider known events. Since knowledge is monotonic and $\mathcal{M}_i$ is a model of total provability, we must have $F \in \mathbf{M}_i(s)$ for every s and every F such that $\Pi_i(s) \subseteq F$. Hence, it follows that $\mathcal{M}_i$ satisfies (M).

To see that 1. and 3. are equivalent, notice that $\Pi_i(s) \in \mathcal{A}_{\Pi_i(s)}$ for every s. If $\mathcal{A}_{\Pi_i}$ is closed under supersets, it follows that $F \in \mathcal{A}_{\Pi_i}(s)$ for every s and every F such that $\Pi(s) \subseteq F$. Thus, $\mathcal{M}_i$ is a model of total provability. To prove that 1. implies 3., notice that if $\mathcal{M}_i$ is a model of total provability, then for every s and every E such that $\Pi_i(s) \subseteq E$, there exists some evidence structure $\mathbf{M}_i$ such that $E \in \mathbf{M}_i(s)$. Thus, by Proposition 1, for every s and every E such that $\Pi_i(s) \subseteq E$ we have $E \in \mathcal{A}_{\Pi_i}(s)$. Since the set of events that i knows at any s is closed under supersets, $\mathcal{A}_{\Pi_i}(s)$ is closed under supersets for every s. □

Notice that Proposition 1 is only needed to prove that 1. implies 3. While total provability is a demanding condition, it has been assumed in some contributions to the literature (see, e.g., [17], [7] and [10]). The following result indicates that total provability models are only possible in a "small" class of models:

Corollary 5. *If $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is a model of total provability, then Π_i is either the finest or the coarsest partition.*

Proof. We show this by contradiction. Suppose that the partition is neither the finest nor the coarsest. The latter implies that there are at least two states, s and s', such that $\Pi_i(s') \cap \Pi_i(s) = \emptyset$. In turn, the former implies that, without loss of generality, we can take either $\Pi_i(s)$ or $\Pi_i(s')$ to be non-singleton.[6] Assume that $|\Pi_i(s)| > 1$.

$$\begin{aligned}
&\Rightarrow \quad \Pi_i(s') \subseteq \Pi_i(s') \cup \{s\} && (1)\\
&\Rightarrow \quad \Pi_i(s') \cup \{s\} \in \mathcal{A}_{\Pi_i}(s') \quad \text{by Proposition 1.} && (2)\\
&\Rightarrow \quad \Pi_i(s') \cup \{s\} \in \mathcal{A}_{\Pi_i}(s) \quad \text{by consistency.} && (3)\\
&\Rightarrow \quad \Pi_i(s) \subseteq \Pi_i(s') \cup \{s\} \quad \text{by measurability.} && (4)\\
&\Rightarrow \quad \Pi_i(s) = \{s\} \quad \text{by } \Pi_i(s) \cap \Pi_i(s') = \emptyset. && (5)
\end{aligned}$$

Since $\Pi_i(s)$ is non-singleton by hypothesis, we reach a contradiction. □

At a fundamental level, Corollary 5 highlights a trade-off between the monotonicity of knowledge and the consistency of evidence. Indeed, by requiring that an agent's knowledge and evidence coincide, total provability requires evidence to be monotonic. However, Corollary 5 shows that the set of provable events, or equivalently the algebra generated by a partition, is only closed under supersets when the agent is either informed about the state or when she knows nothing. Put another way, an agent's knowledge is in general not provable because knowledge is monotonic but evident knowledge is not.[7] The following example illustrates that the conflict between the consistent and monotonic nature of an agent's evidence disappears when the agent is informed about the actual state.[8]

Example 1. *Let $S = \{s_1, s_2, s_3\}$, $\mathcal{E}_i = \{e_1, ..., e_7\}$, and suppose that $\Pi_i(s_j) = \{s_j\}$ for every $j = 1, 2, 3$. Assume, in addition, that messages are distributed across states as follows: $M_i(s_1) = \{e_1, e_2, e_3, e_4\}$, $M_i(s_2) = \{e_2, e_4, e_5, e_6\}$, and $M_i(s_3) = \{e_3, e_4, e_6, e_7\}$. It easy to check that at every state the agent can prove every true event. Thus, $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is a model of total provability.*

It is not hard to see that a model is always of total provability when the partition is the coarsest one. Yet of course, that is not necessarily the case when the partition is the finest one. Thus, the converse of Proposition 5 is false. Figure 1 below illustrates the conflict between monotonicity and consistency by describing a situation in which the agent is not informed about the state but holds nontrivial information.

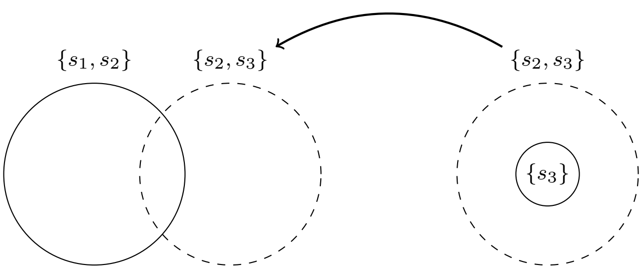

Figure 1: The conflict between consistency and monotonicity.

Figure 1 describes a situation where the agent is informed of s_3 but cannot distinguish between states s_1 and s_2. To see why total provability is impossible in this case, notice that total provability requires the

[6] Indeed, if both are singletons, the fact that the partition is not the finest one implies that there is some state $\bar{s}$ such that $\Pi_i(\bar{s})$ is non-singleton. Hence, $\Pi_i(\bar{s}) \cap \Pi_i(s) = \emptyset$. The argument that follows could then be carried out with $\bar{s}$ and s.

[7] Notice that the conclusion of Corollary 5 can be generalized: when an agent's information is neither the finest nor the coarsest, there are events that are known but not provable or events that can be proven but not be known.

[8] Corollary 5 shows that the conflict also disappears when the agent's partition is the trivial one.

agent to be able to prove both $\{s_3\}$ and $\{s_2, s_3\}$. Indeed, both of these events are known at s_3. But then, consistency implies that the agent must have a message at s_2 that proves the event $\{s_2, s_3\}$. Yet, this is impossible since $\{s_2, s_3\}$ is not known at s_2. Intuitively, total provability entails that the monotonic nature of the agent's knowledge is "inherited" by her proof ability. What Proposition 5 highlights is that, when the partition is neither the finest nor the coarsest, that inherited ability *always* conflicts with consistency (or measurability).

5 Partial provability

Proposition 5 implies that, if an agent's partition is neither the finest nor the coarsest one, she will know events that she cannot prove, even in principle. One might then naturally ask about an agent's proof ability when she is uninformed about the state but still holds non-trivial information. In other words, when is an agent, whose partition is neither the finest nor the coarsest one, able to prove some, *but not every*, event that she knows? Proposition 5 implies that either (M) or (IC) must be relaxed, but in this section we will argue that the answer depends heavily on (M).

Models in which an agent is able to prove some of what she knows can be defined as the complement of those in which she has no evidence:

Definition 5. *$\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is a cheap-talk model if for every s and s', we have that $\mathbf{M}_i(s') = \mathbf{M}_i(s)$.*

Then:

Definition 6. *$\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is a model of* ***partial provability*** *if $\mathcal{M}_i$ is not a cheap-talk model.*

The following property is a joint condition on an agent's knowledge and proof ability:

Definition 7. *$\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is* ***size-measurable*** *(SM) if, for every s, s':*

$$|\mathbf{M}_i(s)| > |\mathbf{M}_i(s')| \Rightarrow |\Pi_i(s')| \geq |\Pi_i(s)|.$$

Intuitively, (SM) requires that the number of events that an agent knows weakly increases with the number of events that she can prove. It is not hard to see that (SM) is a necessary condition for total provability. However, (SM) is independent of both (IC) and (M). The following is the main result of this section:

Proposition 3. *Consider a model $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ in which Π_i is not the finest partition. Then $\mathcal{M}_i$ is a model of partial provability if and only if it fails to satisfy either (M) or (SM).*[9]

The proof of this statement requires a couple of auxiliary results:

Claim 1. *For every s such that $\Pi_i(s)$ is not a singleton,*

$$E \in \mathbf{M}_i(s') \Rightarrow \Pi_i(s) \subseteq E \text{ for every } s' \text{ and every } E. \tag{6}$$

Proof. Take any s such that $\Pi_i(s)$ is not a singleton and assume, contrary to hypothesis, that there is some state s' and event E such that $E \in \mathbf{M}_i(s')$ but $\Pi_i(s) \not\subseteq E$. The latter implies that there is some $\bar{s} \in \Pi_i(s)$ such that $\bar{s} \notin E$. In fact, we must have $\Pi_i(s) \cap E = \emptyset$. To see this, suppose that there is some $\hat{s} \in \Pi_i(s)$ such that $\hat{s} \in E$. Since $E \in \mathbf{M}_i(s')$, consistency entails that $E \in \mathbf{M}_i(\hat{s})$. Thus, $\Pi_i(\hat{s}) \subseteq E$. But since $\hat{s} \in \Pi_i(s)$, it follows that $\Pi_i(s) \subseteq E$. Contradiction.

Since $E \in \mathbf{M}_i(s')$, monotonicity implies that $E \cup \{\bar{s}\} \in \mathbf{M}_i(s')$ for every $\hat{s} \in \Pi(s)$. By consistency, we must then have that $E \cup \{\bar{s}\} \in \mathbf{M}_i(\hat{s})$ for every $\hat{s} \in \Pi_i(s)$. Thus, it follows that $\Pi_i(s) \subseteq E \cup \{s\}$. But since $\Pi_i(s) \cap E = \emptyset$, it must follow that $\Pi_i(s) = \{s\}$. Contradiction. □

Claim 1 entails the following corollary:

Claim 2. *For every s such that $\Pi_i(s)$ is not a singleton,*

$$\mathbf{M}_i(s') \subseteq \mathbf{M}_i(s) \text{ for every } s'. \tag{7}$$

[9] If Π_i is the coarsest partition, $\mathcal{M}_i$ satisfies trivially both (M) and (SM).

Proof. To see this, take any $E \in \mathbf{M}_i(s')$. By Claim 1, we must have that $\Pi_i(s) \subseteq E$. Hence, $s \in E$. Thus, consistency implies that $E \in \mathbf{M}_i(s)$. □

Then, we have:

Proof of Proposition 3: It is easy to see that the *only-if* part holds since both (M) and (SM) are trivially satisfied in cheap-talk models.

To see the *if* part, notice that, by hypothesis, there is a state s at which $\Pi_i(s)$ is not a singleton. Then, (SM) implies that for every s such that $\Pi_i(s)$ is not a singleton we must actually have that $\mathbf{M}_i(s') = \mathbf{M}_i(s)$ for every s'. Suppose that this is not the case. Then, for some states s' and s such that $|\Pi_i(s)| > 1$, $\mathbf{M}_i(s') \subsetneq \mathbf{M}_i(s)$. It follows that $|\mathbf{M}_i(s)| > |\mathbf{M}_i(s')|$. Then, by (SM) we have that $|\Pi_i(s')| \geq |\Pi_i(s)|$. Since $\Pi_i(s)$ is not a singleton, $\Pi_i(s')$ is also not a singleton. But then, Claim 2 implies that $\mathbf{M}_i(s) \subseteq \mathbf{M}_i(s')$, a contradiction. Thus, whenever Π_i is not the finest partition it follows that the model is a cheap-talk model. □

We finish the paper by stating an important corollary of Claim 2:

Corollary 6. *Fix any model $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ in which $\Pi_i(s)$ is non-singleton at every s. Then, $\mathcal{M}_i$ is a model of partial provability if and only if $\mathcal{M}_i$ is not monotonic.*

Corollary 6 shows that, when moving from total to partial provability one has to dispense with (M). In other words, (M) is necessary for total provability when the partition is either the finest or the coarsest one, but necessary for cheap-talk under any other partition.

6 References

[1] Ben-Porath, E. and Lipman, B. L. (2012). Implementation with partial provability. *Journal of Economic Theory* 147(5): 1689–1724.

[2] Ben-Porath, E., Dekel, E. and Lipman, B. L. (2019). Mechanisms with evidence: Commitment and robustness. *Econometrica* 87(2): 529–566.

[3] Binmore, K. and Brandeburger, A. (1988). Common knowledge and game theory. *Working Paper: Department of Economics, University of Michigan.*

[4] Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica* 50(6): 1431–1451.

[5] Forges, F. and Koessler, F. (2005). Communication equilibria with partially verifiable types. *Journal of Mathematical Economics* 41(7): 793–811.

[6] Okuno-Fujiwara, M., Postlewaite, A. and Suzumura, K. (1990). Strategic information revelation. *Review of Economic Studies* 57(1): 25–47.

[7] Giovannoni, F. and Seidmann, D. J. (2007). Secrecy, two-sided bias and the value of evidence. *Games and Economic Behavior* 59(2): 296-315.

[8] Green, J. R. and Laffont, J. J. (1986). Partially verifiable information and mechanism design. *Review of Economic Studies* 53(3): 447–456.

[9] Hagenbach, J., Koessler, F. and Perez-Richet, E. (2014). Certifiable Pre-Play Communication: Full Disclosure. *Econometrica* 82(3): 1093–1131.

[10] Hagenbach, J. and Koessler, F. (2017). Simple versus rich language in disclosure games. *Review of Economic Design* 21(3): 163–175.

[11] Hart, S., Kremer, I. and Perry, M. (2017). Evidence games: Truth and commitment. *American Economic Review* 107(3): 690–713.

[12] Kartik, N. and Tercieux, O. (2012). Implementation with evidence. *Theoretical Economics* 7(2): 323–355.

[13] Koessler, F. (2004). Strategic knowledge sharing in Bayesian games. *Games and Economic Behavior* 48(2): 292–320.

[14] Lipman, B. L. and Seppi, D. J. (1995). Robust inference in communication games with partial provability. *Journal of Economic Theory* 66(2): 370–405.

[15] Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics* 12(2): 380–391.

[16] Milgrom, P. (1981). An axiomatic characterization of common knowledge. *Econometrica* 49(1): 219–222.

[17] Milgrom, P. and Roberts, J. (1986). Relying on the information of interested parties. *The RAND Journal of Economics* 17(1): 18–32.