Provability: from evidence to knowledge

Esteban Peralta* and Fernando Tohmé†

Abstract

This note explores 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 varies with knowledge. Fixing a partition, an event is said to be provable if there exists a message signaling it in some evidence structure. We show that an event is provable if and only if it is self-evident—i.e., known at everyone of its states. We also find that knowledge is provable only if the partition of the state space is either the coarsest or the finest one. A relaxation of the conditions for provability, requiring only that some known events are provable, allows for other partitions as well. But this is only possible if we dispense with the property of monotonicity of provability. These results offers a novel foundation for knowledge and common knowledge, in which they emerge from the evidence held by the agents.

Keywords: Knowledge - Evidence - Common Knowledge - Implementation.

1 Introduction

Evidence can solve coordination problems by creating common knowledge. Suppose that two agents privately observe both a signal and a set of messages that varies with the signal they receive. Each agent is deciding whether to invest on a project that is only successful if the economic conditions are "healthy". Each agent decides to join if and only if it is a commonly known fact among them that the economic conditions are good. In this paper we study how the information conveyed by the private signals received by the agents can become common knowledge. A rather surprising result is that in the context of the coordination problem, the agents will have common knowledge of the health of the economic situation if and only if *both have evidence* that the economic conditions are good.

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 the settings in which messages convey no information and only cheap-talk communication is possible [4]. Since an agent's ability to offer evidence limits her ability to misrepresent her information, the study of models with partial provability has attracted a considerable amount of attention both in communication settings and mechanism design problems.[1]

Despite its importance, how an agent's ability to prove depends on her information is not yet well understood. This is somewhat surprising as the literature has spanned the whole range of proof abilities, from cheap talk communication to total provability; namely, situations where an agent can prove every piece of her knowledge (see, e.g., [17], [7] and [10]). This note intends to uncover the relationship between an agent's evidence and knowledge in a model where 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, agents' are not necessarily informed about the actual state but they know every event that they can prove.

The starting point of this paper is the observation that imposing this measurability requirement implies that the class of events that are provable with respect to a given partition is given by the algebra generated by

*Department of Economics and Stephen M. Ross School of Business, University of Michigan. I am grateful to Shaowei Ke and Larry Samuelson for very helpful discussions. The usual disclaimer applies.

†Departamento de Economía, Universidad Nacional del Sur and CONICET

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

the partition—minus the empty set. This observation has several consequences.[2]

We show that an event is *provable* if and only if it is self-evident and *known*. Moreover, we show that an event is public [16]—i.e. commonly known at every one of its states—if and only if it is provable by every agent. Thus, **an event is commonly known if and only if it is implied by a mutually provable event**.

Since the algebra generated by a partition is the richest evidence structure that is measurable with respect to the partition [13], these results offer more than a novel characterization of knowledge and self-evident knowledge. They can be interpreted as explaining *when*, or *why*, an event is known, evidently or not. Indeed, the algebra generated by a partition can be interpreted as describing the messages that exist at every state, regardless of whether some of these messages are accessible or not to a given agent. Thus, the main philosophical conclusion that can be derived from our results is that knowledge *emerges* from evidence. Therefore, evidence should be taken as the primitive notion in epistemic analyses.

On the other hand, our results indicate that the conditions for total provability are too demanding. Total provability is only achieved in models in which the partition of states is either the coarsest or the finest one. To see which features are responsible for the surprising identity between common and mutual knowledge, we relax the framework by considering models of partial provability. In these models, in which cheap-talk is not possible, we obtain positive results for a wider class of partitions. But this is only at the cost of dropping the condition of monotonicity of provability. Thus, we find that an agent is able to prove every event she knows, only a the top and bottom in the lattice of partitions, and that if evidence is monotonic. Otherwise, there will be known events that she will not able to prove.

The plan of the paper is as follows. Section 2 formally presents the notions of knowledge, evidence and proof, as well as the fundamental condition of measurability of evidence structures with respect to information partitions. Section 3 introduces the main claims in the paper, indicating new relations among the aforementioned concepts. Section 4 revisits the results found in the previous section, showing that they amount to say that the models are of total provability. In section 5 we relax this condition to partial provability, finding a wider class of partitions in which some known events can be proven. Section 6 concludes.

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 a partition of S for i. Thus, $\Pi_i(s)$ is the collection of states that agent i considers possible when the state is s. An *information structure* 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 $K_I^{m+1}(E) = K_I(K_I^m(E))$, where $K_I^1(E) = K_I(E)$. Thus, define the event E is common knowledge, $CK(E)$, as follows:

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

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

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 **evident knowledge for** i if i knows E at every one of its states; i.e., if $E \subseteq K_i(E)$. Since $K_i(E) \subseteq E$ for every i and every E, one can equivalently state that an event E is evident knowledge for i if $K_i(E) = E$. Notice that any event that is evident knowledge 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 evident knowledge 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** if it is evident knowledge for every i; i.e., if $E \in \mathcal{A}_{\Pi_i}$ 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.:

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

Define $\mathbf{M}_i := \bigcup_{s \in S} \mathbf{M}_i(s)$ to be an **evidence structure for agent** i. Define $\mathbf{M} := \{\mathbf{M}_i\}_{i \in A}$.

By construction, evidence structures satisfy two properties. First, for any s and any $E \in \mathbf{M}_i(s)$, we must 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 $\mathbf{M}_i(s) = \mathbf{M}_i(s')$. A **model** for I is the class $\mathcal{M} = \{\mathcal{M}_i\}_{i \in I}$ where, for every i, $\mathcal{M}_i$ is a model for i.

We have that:

Proposition 1. *In a model $\mathcal{M}$, each agent i knows every event that she can prove.*

Proof. Take any two states s and s' such that $\Pi(s) = \Pi(s')$ and any event $E \in \mathbf{M}_i(s)$. Then, $\Pi(s) \subseteq E$. Hence, $s' \in E$. Thus, consistency implies that $E \in M_i(s')$. Hence, $\mathbf{M}_i(s) \subseteq \mathbf{M}_i(s')$. The same argument applies to any $E \in \mathbf{M}_i(s')$, so that $\mathbf{M}_i(s') \subseteq \mathbf{M}_i(s)$.

To see the converse, assume that $\mathbf{M}_i$ is measurable with respect to Π but that there is an event $E \in \mathbf{M}_i(s)$ such that $\Pi(s) \not\subseteq E$. Since provable events must be true, we must have $s \in E$. Hence, there is a state $\bar{s} \in \Pi(s)$ such that both $\bar{s} \neq s$ and $\bar{s} \notin E$. It must follow that $E \notin \mathbf{M}_i(\bar{s})$. Yet, $\Pi(s) = \Pi(\bar{s})$, contradicting that $\mathbf{M}_i$ is measurable with respect to Π. □

[3] Alternatively, common knowledge can be described in terms of the meet of Π; namely, the finest common coarsening of Π.
[4] Thus, public events can be written as the union of elements of the meet of Π.

The definition of model extends the meaning of verifiable information to situations in which agents are not necessarily informed about the state at which they are [13]. Indeed, a piece of information is verifiable precisely when it varies, not only with the state, but with the agent's information.

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 Provable events

The next result states that an agent can only offer evidence of what is self-evident to her, regardless of her information.[5]

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

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 $\Pi_i(s') \not\subseteq E$. Yet consistency entails that $E \in \mathbf{M}_i(s')$, contradicting that $\mathbf{M}_i$ is measurable with respect to Π_i. □

Intuitively, Lemma 1 shows that consistency strengthens the requirement that an agent's evidence must be measurable with respect to her knowledge; an agent is not only restricted to offer evidence of what she knows, but can only offer evidence of what it is self-evident to her. Thus, $\mathcal{A}_{\Pi_i}$ can be seen as a "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 *can be* proved—i.e., are proved in *some* measurable evidence structure—coincides with the set of self-evident events generated by the partition.

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. Thus, i might not have evidence for a provable event in some (measurable) evidence structures. 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 Lemma 1, it follows that $E \in \mathcal{A}_{\Pi_i}(s)$. □

This result strengthens Lemma 1 in the sense that it asserts that an agent can prove *only those* events that are self-evident for her. Thus, an agent's self-evident knowledge describes the scope, or "upper bound", of her proof ability.

3.1 Provable knowledge

It is well known that an agent's knowledge can be characterized in terms of her evident knowledge (see, e.g., [3]). Thus, an immediate implication of Proposition 2 is that an agent's knowledge can be characterized in terms of what is provable by her.

Corollary 1. *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.

The next result indicates that monotonicity of evidence is equivalent the provability of known events.

Proposition 3. *Agent* i*'s knowledge is provable if and only if* $\mathcal{A}_{\Pi_i}(s)$ *is closed under supersets for every* 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 i) he didn't highlight why this is the case (consistency) and ii) he didn't study the consequences of this property.

Proof. Fix any agent i. To see the *if* part, 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, i's knowledge is provable.

For the *only if* part, notice that if i's knowledge is provable, 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 Lemma 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. □

The following result states the limits to the provability of known events:

Proposition 4. *Agent i's knowledge is provable only if Π_i is either the finest or the coarsest partition.*

Proof. We can show this by contradiction. Suppose that the partition is neither the finest nor the coarsest. The latter implies that there are 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 the former.

$$\Rightarrow \quad \Pi_i(s') \subseteq \Pi_i(s') \cup \{s\} \tag{1}$$
$$\Rightarrow \quad \Pi_i(s') \cup \{s\} \in \mathcal{A}_{\Pi_i}(s') \quad \text{by total provability.} \tag{2}$$
$$\Rightarrow \quad \Pi_i(s') \cup \{s\} \in \mathcal{A}_{\Pi_i}(s) \quad \text{by consistency.} \tag{3}$$
$$\Rightarrow \quad \Pi_i(s) \subseteq \Pi_i(s') \cup \{s\} \quad \text{by measurability.} \tag{4}$$
$$\Rightarrow \quad \Pi_i(s) = \{s\} \quad \text{by } \Pi_i(s) \cap \Pi_i(s') = \emptyset. \tag{5}$$

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

At a fundamental level, Proposition 4 highlights a trade-off between a property of knowledge and a property of evidence; namely, monotonicity and consistency. Indeed, notice that by requiring that an agent's knowledge and evidence coincide, total provability requires evidence to be monotonic. However, Proposition 4 shows that the set of provable events, or equivalently the algebra generated by a partition, is only closed under supersets when the agent is informed about the state or knows nothing. Put another way, an agent's knowledge is in general not provable because knowledge is monotonic but evident knowledge is not.[7]

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.

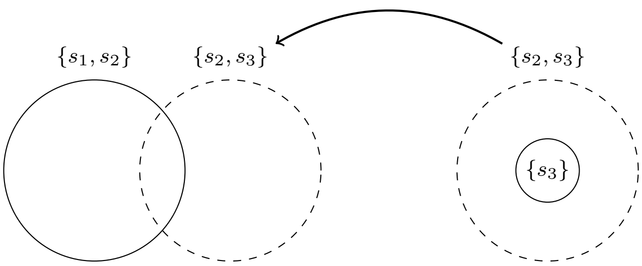

Figure 1: The conflict between consistency and monotonicity.

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 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 posses 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 4 highlights is that, when the

[6]Indeed, if both are singleton, 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 Proposition 4 can be generalized. Namely, when an agent's information is neither the finest nor the coarsest, there are events that are either known but cannot be proved or events that are proved but cannot be known.

partition is neither the finest nor the coarsest, that inherited ability *always* conflicts with consistency (or measurability).

3.2 Provable common knowledge

Since common knowledge can be characterized in terms of the agents' evident knowledge (see, e.g., [3]), Proposition 2 and Corollary 1 readily extend to the class of events that are commonly known by a group of agents. This is the content of the next corollary.

Corollary 2. *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.*

The first statement in this result implies that the class of events that are *not* public are precisely those that *cannot* be proved by *some* agent. Conversely, every event that *can be* proved by all agents must be evident common knowledge among them. On the other hand, the second statement in Corollary 2 offers a foundation for common knowledge in terms of (mutual) evidence.

Corollary 2 not only rationalizes how common knowledge and evident common knowledge can be attained, but also highlights that the infinite regress behind them is captured by the less demanding notion of *mutual* provability.

4 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 with each other. A model in which this is the case can be defined as follows:

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 where the agent's information is evidence-based in the sense that the agent's knowledge is restricted by her proof-ability. Thus, total provability models can be seen as models that capture the idea that *knowledge must be grounded in evidence.*

4.1 Characterization

Since an agent must know every event that she can prove, total provability models are models in which the agent knows an event if and only if she can prove it. Thus, total provability not only asks the agent to be able to prove the signal she receives, but also entails that the monotonicity in the agent's knowledge must be inherited by her proof ability.

Definition 3. *$\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is **information-certifiable** (IC) if $\Pi_i(s) \in \mathbf{M}_i(s)$ for every s.*

In words: information-certifiable models are models in which, at every state, the agent can prove the signal she receives. The condition thus extends the notion of complete provability in [14] to situations where the agent is not necessarily informed about the state. Information-certifiability is closely related to the notion of own-type certifiability proposed by [9].

Definition 4. *$\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is **monotonic** (M) if, for every s:*

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

Notice that (M) is a property of events, not of evidence.

The following result characterizes the class of evidence-based models:

Proposition 5. *$\mathcal{M}_i$ is a model of total provability if and only if $\mathcal{M}_i$ satisfies (M) and (IC).*

Proof. To see the *if* part, 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.

For the *only-if* part, 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). □

Notice that the only-if part uses the assumption that an agent's evidence must be measurable with respect to her information. This is illustrated in the following example.

Example 1. *Let $S = \{s_1, s_2, s_3\}$, $\mathcal{E}_i = \{e_1, e_2\}$, and suppose that $\Pi_i(s_j) = S$ for $j = 1, 2, 3$. Assume, in addition, that messages are distributed across states as follows: $M_i(s_1) = \{e_1, e_2\}$ and $M_i(s_2) = M_i(s_3) = \{e_2\}$. It is not hard to see that this is a model of total provability that satisfies (IC). However, monotonicity fails since $\{s_1\} \in \mathbf{M}_i(s_1)$ but, for example, $\{s_1, s_2\} \notin \mathbf{M}_i(s_1)$.*

4.2 Necessary conditions

While total provability is a demanding condition, it has been assumed by some papers in the literature (see, e.g., [17], [7] and [10]). The following result, which is an immediate consequence of Proposition 4, indicates that total provability models cannot in general be assumed since it is condition that is only valid in very specific cases:

Proposition 6. *If $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is a model of total provability, then $\mathcal{M}_i$ is either the finest or the coarsest partition.*

Proposition 6 highlights again the fact that when a partition is neither the finest nor the coarsest, monotonicity *always* conflicts with consistency. Indeed, the following example illustrates that this conflict between the consistent and monotonic nature of the agent's evidence disappears when the agent is informed about the actual state.[8]

Example 2. *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 is not hard to see that, at every state, the agent can prove every true event. Thus, this 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. But that is not the case when the partition is the finest one, since $S \in \mathbf{M}_i(s)$ for every $s \in S$. Thus, the converse of Proposition 6 is false.

5 Partial provability

In light of Propositions 4 and 6 one might naturally ask about an agent's proof ability when she is uninformed about the state but still holds non-trivial information. Then, what guarantees that an agent will be able to prove some, *but not all*, of the pieces of the knowledge that she holds? Models in which this is possible can be defined in contrast to *cheap-talk* models:

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

Then:

Definition 6. *A model $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ is one 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.

[8] Proposition 6 shows that the conflict also disappears when the agent's information is trivial.

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). Then:

Proposition 7. *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 does not satisfy (M) and (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,*

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

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 7: 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 can conclude, from Proposition 7 that in order to get partial provability we have to relax either (IC) or (M). From Claim 2 we can derive the following important corollary that shows that actually we have to dispense of (M):

Corollary 3. *Consider a model* $\mathcal{M}_i = (\Pi_i, \mathbf{M}_i)$ *in which* Π_i *is nor the coarsest or the finest partition. If* $\Pi_i(s)$ *is non-singleton at every* s*,* $\mathcal{M}_i$ *is a model of partial provability if and only if* $\mathcal{M}_i$ *is not monotonic.*

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

6 Conclusion

In this brief note we have shown a remarkable property of knowledge and in particular of common knowledge. Namely that it can be seen as gained up from evidence provided by signals received by the agents, be it directly or through the exchange of messages.

Furthermore, we saw that mutual provability is enough to yield common knowledge, because the former defines public events that lurk behind the latter. Then, back to our initial example, if each one of the two agents deciding on investing on a risky project gets evidence that the economic context is good, that is enough to make common knowledge the fact that they have evidence of the health of the economy, without having to communicate among them.

These results ensue in frameworks of total provability, only the finest and coarsest partitions. If we relax the requirement to partial provability we find that other partitions also allow proving known events. But that at the cost of dropping the monotonicity of the evidence, which was one of the main properties ensuring the equivalence between common and mutual knowledge.

These results, trivial as their proofs are, provide a full picture of the relations between evidence and knowledge, with some surprising consequences. This indicates that there might exist further room for exploration and understanding on this topic.

7 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.