Information Hierarchies

Benjamin Brooks, Alexander Frankel, and Emir Kamenica*

September 2019

Abstract

If Anne knows more than Bob about the state of the world, she may or may not know what Bob thinks, but it is always possible that she does. In other words, if the distribution of Anne's first-order belief is a mean-preserving spread of the distribution of Bob's first-order belief, we can construct signals for Anne and Bob that induce these distributions of beliefs and provide Anne with full information about Bob's belief. We establish that with more agents, the analogous result does not hold. It might be that Anne knows more than Bob and Charles, who in turn both know more than David, yet what they know about the state precludes the possibility that Anne knows what Bob and Charles think and that everyone knows what David thinks. More generally, we define an *information hierarchy* as a partially ordered set and ask whether higher elements having more information about the state always makes the hierarchy compatible with higher elements knowing the beliefs of lower elements. We show that the answer is affirmative if and only if the graph of the hierarchy is a forest. We discuss applications of this result to rationalizing a decision maker's reaction to unknown sources of information and to information design in hierarchical vs. non-hierarchical organizations.

JEL classification: C70; D82; D83; D85
Keywords: Information, beliefs, signals, Blackwell order, networks

*Brooks: Department of Economics, University of Chicago, babrooks@uchicago.edu. Frankel: Booth School of Business, University of Chicago, afrankel@uchicago.edu. Kamenica: Booth School of Business, University of Chicago, Emir.Kamenica@uchicago.edu. Kamenica thanks the Neubauer Family Faculty Fellowship for research support.

1 Introduction

There are two distinct things we might mean when we say, "Anne is more informed than Bob." One is that Anne's information about some state of the world is more accurate than Bob's. The other is that Anne knows everything that Bob knows. Of course, the latter sense implies the former: if Anne knows everything that Bob knows, her overall information is necessarily more accurate. Moreover, Blackwell's Theorem (1953) tells us that if Anne's information is more accurate, then it is always possible that Anne knows everything that Bob knows, whatever the extent of Anne's or Bob's knowledge about the state might be. In formal terms, if the distribution of Anne's belief is a mean-preserving spread of the distribution of Bob's belief, we can always construct signals for Anne and Bob that induce these distributions of beliefs and ensure that Anne knows everything that Bob knows.

In this paper, we explore the relationship between these two notions of "more informed" more generally. Perhaps surprisingly, the aforementioned implication of Blackwell's Theorem does not extend beyond three players. We construct an example where Anne's information is more accurate than Bob's and Charles's, whose information is in turn more accurate than David's, and yet it cannot be that Anne knows everything that Bob and Charles know and that all three know everything that David knows. Formally, we construct four distributions of posteriors, τ_A, τ_B, τ_C, and τ_D such that τ_A is a mean-preserving spread[1] of τ_B and τ_C, which are in turn mean-preserving spreads of τ_D, and show there do not exist four signals π_A, π_B, π_C, π_D that induce these four distributions of posteriors and have the property that observing π_A suffices to know the realizations of π_B, π_C, and π_D, and observing π_B or π_C suffices to know the realization of π_D. It is always the case that a less-informed person cannot know what the more-informed people think; in this example, a more-informed person cannot know what less-informed people think.

To examine this issue in full generality, we introduce the notion of an *information hierarchy*. An information hierarchy is simply a partially ordered set. Each information hierarchy is associated with an undirected graph, whose nodes are the elements of the set and whose edges are determined by the partial order.[2] We consider allocations of distributions of beliefs to the elements of the

[1] Recall that a more informative signal (in the sense of Blackwell 1953) induces a more dispersed distribution of posterior beliefs.

[2] The standard representation of a partially ordered set as a directed graph encodes the partial order by placing

hierarchy that are *monotone*, in the sense that higher elements have more accurate information. Given such an allocation, we may ask whether those beliefs could be induced by monotonically allocating a signal to each element, in that higher elements can infer the signal realizations of lower elements. If every monotone belief allocation can be induced by a monotone signal allocation, then we say that the information hierarchy is *universally constructible*. If a hierarchy is universally constructible, then the Blackwell ranking exactly captures implications for beliefs of the requirement that higher elements know the signals of lower elements.

Our main theorem establishes that an information hierarchy is universally constructible if and only if its undirected graph is a forest, i.e., there is at most one path between any two elements. For example, the aforementioned four-person example entails an information hierarchy whose graph is not a forest: there are two paths from Anne to David, one through Bob and one through Charles. In contrast, the undirected graph in the two-person example, from the first paragraph, is a forest.[3] This is precisely why in the four-person but not in the two-person example, it was possible for the extent of knowledge about the state to preclude an individual with more accurate information from knowing everything known by the less informed.

The proof of the "if" direction of our theorem is relatively straightforward, though it does require establishing a novel information-theoretic result that might be of independent interest. Under the hypothesis that the undirected graph is a forest, we provide an algorithm that iteratively constructs a monotone signal allocation inducing any given monotone belief allocation. The "only if" direction is considerably more involved. The proof relies on three key ideas: First, we show that a hierarchy is universally constructible only if its *closed subhierarchies* are also universally constructible. In this context, "closed" means that all of the elements in the hierarchy that are between two elements of the subhierarchy are also in the subhierarchy. Second, we show that any hierarchy that is not a forest must contain a closed subhierarchy taking one of two forms: either its undirected graph is a *crown* or it is a *union of undirected paths*. The latter may be seen as a generalization of the four-person example described above. Finally, we present monotone belief allocations for each of these subhierarchy forms that we show cannot be induced by monotone signal allocations.

an edge from n to n' if n covers n', i.e., if $n > n'$ and there is no n'' such that $n > n'' > n'$. We associate with each information hierarchy the undirected version of this directed graph.

[3] Indeed, the graph of any information hierarchy with two or three elements is necessarily a forest.

Our work contributes to the growing literature on information design (Kamenica and Gentzkow, 2011; Bergemann and Morris, 2016). In particular, we study what forms of distributed information can be implemented via signals, when there are constraints that require some signals to be refined by others. More broadly, we explore the relationship between different orderings on information and experiments (Blackwell, 1951, 1953). Our contribution extends this theory to the case where many sources of information are being compared simultaneously. Relatedly, Bergemann and Morris (2016) study an extension of the Blackwell ranking to type spaces, which can be understood as a ranking on profiles of signals. We also study how signals can be combined to produce more informative signals. Gentzkow and Kamenica (2017) study this issue in the context of a communication game with one receiver who combines information from signals provided by multiple senders. Börgers, Hernando-Veciana and Krähmer (2013) study the interaction between signals from the perspective of whether signals are substitutes or complements for one another. Our inquiry also leads us to a pure graph-theoretic question of whether a partially ordered set contains subsets of a particular form. This subject has been studied in combinatorics and graph theory (e.g., Lu, 2014) and is applied by Curello and Sinander (2019) to study rankings on preference relations, although our inquiry is distinct in that we restrict attention to closed subhierarchies.

The rest of the paper proceeds as follows. Section 2 describes our model of information hierarchies. Section 3 presents several examples of hierarchies, and discusses which of them are universally constructible and which are not. Section 4 presents our main result and a sketch of the proof. Section 5 discusses two applications of our results: (i) rationalizability of reactions to unknown sources of information, and (ii) information design in hierarchical vs. non-hierarchical organizations. Section 6 briefly concludes. All omitted proofs are in the Appendix.

2 Set-up

2.1 Information hierarchy

There is an *information hierarchy* H, which is a finite partially ordered set $(N, \geq)$ with the corresponding strict order $>$. Since we will heavily rely on graph-theoretic representations of $(N, \geq)$, we refer to elements of N as *nodes*. Nodes n and n' are *comparable* if $n \geq n'$ or $n' \geq n$. Given

$n, n' \in N$, n *covers* n' if $n > n'$ and there does not exist $n'' \in N$ with $n > n'' > n'$.

A subset of nodes $N' \subseteq N$ induces the *subhierarchy* $H' = (N', \geq)$, with the partial order being the restriction of $\geq$ on N'. The *directed graph* of H', denoted $G(H')$, is the pair (N', E), where N' is the set of nodes, $E \subseteq N' \times N'$ is the set of directed edges, and $(n, n') \in E$ if n covers n'. A *directed path* from n to n' in (N', E) is an alternating sequence of vertices and directed edges $(n_0, e_0, ...n_{L-1}, e_{L-1}, n_L)$, where $L > 0$, $n_0 = n$, $n_L = n'$, $n_l \in N'$ for all $l \in \{0, ..., L\}$, $e_l = (n_l, n_{l+1}) \in E$ for all $l \in \{0, ..., L-1\}$, and $l \neq l' \Rightarrow e_l \neq e_{l'}$. The *undirected graph* of H', denoted $\tilde{G}(H')$, is the pair $\left(N', \tilde{E}\right)$, where N' is the set of nodes, $\tilde{E} \subseteq \{\tilde{e} \subseteq N' | \, |\tilde{e}| = 2\}$ is the set of undirected edges, and $\{n, n'\} \in \tilde{E}$ if n covers n' or n' covers n. An *undirected path* from n to n' in $\left(N', \tilde{E}\right)$ is an alternating sequence of vertices and undirected edges $(n_0, \tilde{e}_0, ...n_{L-1}, \tilde{e}_{L-1}, n_L)$, where $L > 0$, $n_0 = n$, $n_L = n'$, $n_l \in N'$ for all $l \in \{0, ..., L\}$, $\tilde{e}_l = \{n_l, n_{l+1}\} \in \tilde{E}$ for all $l \in \{0, ..., L-1\}$, and $l \neq l' \Rightarrow \tilde{e}_l \neq \tilde{e}_{l'}$. A *cycle* in $\left(N', \tilde{E}\right)$ is an undirected path from n to n. We say H' is *cyclic* if $\tilde{G}(H')$ contains a cycle.

Note that the directed and undirected graphs of $H' = (N', \geq)$ are defined with respect to the covering relation in N, not in N'. That is, suppose $n, n' \in N'$ and $n > n'$. Then, $G(H')$ and $\tilde{G}(H')$ have an edge between n and n' only if there is no $n'' \in N$ with $n > n'' > n'$; absence of such n'' in N' does not suffice to generate an edge. The graphs depict what is often termed the transitive reduction: if $n > n'$, there is a path from n to n', but there is an edge from n to n' only if there is no node between them.

An undirected graph is a *tree* if there is exactly one undirected path between any two nodes. An undirected graph is a *forest* if there is at most one undirected path between any two nodes, i.e., there are no cycles. Thus, a forest is a union of disjoint trees.

2.2 States and signals

Given a finite state space Ω and a prior $\mu_0 \in \Delta(\Omega)$, a *signal* π is a finite partition of $\Omega \times [0, 1]$ s.t. $\pi \subset S$, where S is the set of non-empty Lebesgue-measurable subsets of $\Omega \times [0, 1]$ (Green and Stokey, 1978; Gentzkow and Kamenica, 2017). An element $s \in S$ is a *signal realization*.

The interpretation of this formalism is that a random variable x drawn uniformly from $[0, 1]$ determines the signal realization conditional on the state. Let $p(s|\omega) = \lambda(\{x| \, (\omega, x) \in s\})$ and

$p(s) = \sum_{\omega \in \Omega} p(s|\omega) \mu_0(\omega)$, where $\lambda(\cdot)$ denotes the Lebesgue measure. That is, $p(s|\omega)$ is the conditional probability of s given ω and $p(s)$ is the unconditional probability of s. We denote the set of all signals by Π.

We denote the refinement order on Π by $\trianglerighteq$, i.e., given $\pi, \pi' \in \Pi$, we write $\pi \trianglerighteq \pi'$ if every element of π is a subset of some element of π'. The pair $(\Pi, \trianglerighteq)$ is a lattice and we let $\vee$ denote the join, i.e., $\pi \vee \pi'$ is the coarsest refinement of both π and π'. For any set P, we denote the join of all its elements by $\bigvee P$.

2.3 Distributions of posteriors

Given Ω and μ_0, a *distribution of beliefs*, denoted by τ, is an element of $\Delta\Delta\Omega$ that has finite support and satisfies $\mathbb{E}_\tau[\mu] = \mu_0$. We partially order distributions of beliefs by informativeness in the sense of Blackwell (1951) and write $\tau \succsim \tau'$ if τ is a mean-preserving spread of τ'. We let $\overline{\tau}$ denote the maximally informative distribution of beliefs (whose support contains only degenerate beliefs) and $\underline{\tau}$ denote the minimally informative distribution of beliefs (that puts probability one on the prior).

Observing a signal realization s s.t. $p(s) > 0$ generates a unique posterior belief, where the probability of ω given s is[4]

$$\mu_s(\omega) = \frac{p(s|\omega)\mu_0(\omega)}{p(s)}.$$

Each signal π induces a distribution of posteriors denoted by $\langle\pi\rangle$, according to

$$\langle\pi\rangle(\mu) = \sum_{\{s \in \pi | \mu_s = \mu\}} p(s).$$

Note that $\pi \trianglerighteq \pi' \Rightarrow \langle\pi\rangle \succsim \langle\pi'\rangle$.

Given signal π, let $\tilde{\mu}_\pi$ denote the belief-valued random variable that reflects the posterior induced by the observation of the signal realization from π.

[4] For those s with $p(s) = 0$, set μ_s to be an arbitrary belief.

2.4 Beliefs and signals in hierarchies

A *belief allocation* on H under (Ω, μ_0) is a map that assigns a distribution of beliefs to every node in N. A belief allocation β is *monotone* (with respect to H) if $n \geq n' \Rightarrow \beta(n) \succsim \beta(n')$, i.e., if greater nodes are Blackwell more informed than lower nodes.

A *signal allocation* on H under (Ω, μ_0) is a map that assigns a signal to every node in N. We say that a signal allocation σ is *monotone* (with respect to H) if $n \geq n' \Rightarrow \sigma(n) \trianglerighteq \sigma(n')$. In other words, a signal allocation specifies a signal for every node, and it is monotone if greater nodes have signals that refine lower nodes' signals. A signal allocation σ *induces* a belief allocation β if for all n, $\beta(n) = \langle \sigma(n) \rangle$.

2.5 Universal constructibility

We say the information hierarchy H is *K-universally constructible* if for every (Ω, μ_0) s.t. $|\Omega| = K$, every monotone belief allocation on H under (Ω, μ_0) is induced by some monotone signal allocation on H under (Ω, μ_0). An immediate implication of this definition is that if a hierarchy is K-universally constructible, then it is K'-universally constructible for all $K' \leq K$. We say a hierarchy is *universally constructible* if it is K-universally constructible for all $K \in \mathbb{N}$. We note for future reference that the notions of K-universal constructibility and universal constructibility also apply to subhierarchies, in the obvious manner.

Note that universal constructibility is required to hold across all priors (given the cardinality of the state space). As shall be clear from the proof, all of our results would go through if we fixed any particular full-support prior. Indeed, we have required universal constructibility to hold across all priors in order to make it clear that our results do not depend on the choice of a particular prior. It is even possible to reformulate our theorem in a "prior-free" manner, and replace distributions of beliefs with experiments as in Blackwell (1951, 1953), i.e., a pair of an abstract signal space and conditional distributions over signals given the state. By explicitly modeling a prior, we can normalize agents' signals to be beliefs, which simplifies our exposition.

3 Examples

In this section, we present several examples of information hierarchies and discuss which of them are universally constructible. Along the way, we establish some Lemmas and intuitions that will play a central role in the proof of our main result.

3.1 Examples of information hierarchies

Example 1 (Chain)**.** There are four individuals—A, B, C, and D—ranked in alphabetical order: A ranks above B, who ranks above C, who ranks above D. The set of nodes is $N = \{A, B, C, D\}$, and $\geq$ reflects the ranking relation (i.e., $A \geq B \geq C \geq D$). Figure 1a depicts the directed graph of the information hierarchy $(N, \geq)$.

Example 2 (Tree)**.** There is a small organization that consists of a president (P) who has two deputies ($D1$ and $D2$), each of whom has two assistants ($A1$ and $A2$; $A3$ and $A4$). The partial order reflects the organizational hierarchy, with $P \geq D1$, $D2$; $D1 \geq A1, A2$; and $D2 \geq A3, A4$. Figure 1b depicts the directed graph.

Example 3 (Diamond)**.** Returning to the example from the introduction, there is an organization whose president (A) has two deputies (B and C) that share an assistant (D). We refer to this hierarchy—with $A \geq B, C$ and $B, C \geq D$—as the *diamond*, depicted in Figure 1c.

Example 4 (Set-inclusion)**.** Elements of an information hierarchy need not represent individuals. Suppose there are three sources of information: X, Y, and Z. A decision maker has access to some (or all) of these sources. The nodes in this *set-inclusion* hierarchy are the possible collections of the information sources, i.e., $N = \{X, Y, Z, XY, XZ, YZ, XYZ\}$, and $\geq$ denotes the inclusion order: $XYZ \geq XY, XZ, YZ$; $XY \geq X, Y$; $XZ \geq X, Z$; $YZ \geq Y, Z$. Figure 1d depicts the directed graph of the hierarchy.

Example 5 (Crown)**.** There are two parents, F and M, who have two children, S and D. We refer to this information hierarchy, with $\geq$ reflecting the parenting relation ($F \geq S, D$; $M \geq S, D$), as the *crown*, depicted in Figure 1e.[5]

[5] More specifically, this hierarchy is a 4-crown. More generally, the n-crown is defined as a partially ordered set in which half of the nodes are maximal and half are minimal, and there is a single cycle that contains all of the nodes.

Figure 1: Example Hierarchies

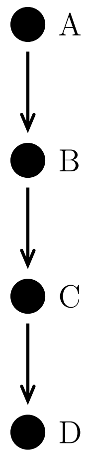

(a) Chain

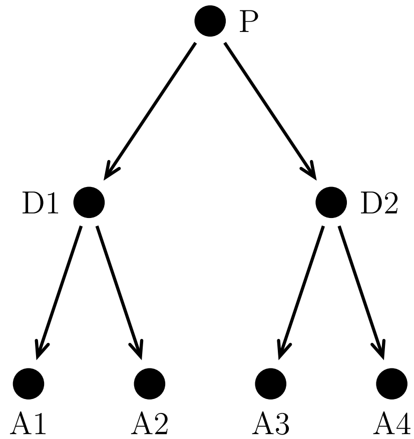

(b) Tree

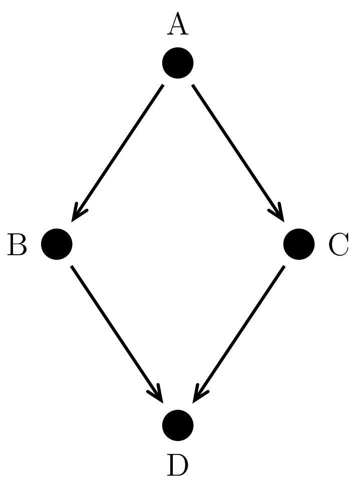

(c) Diamond

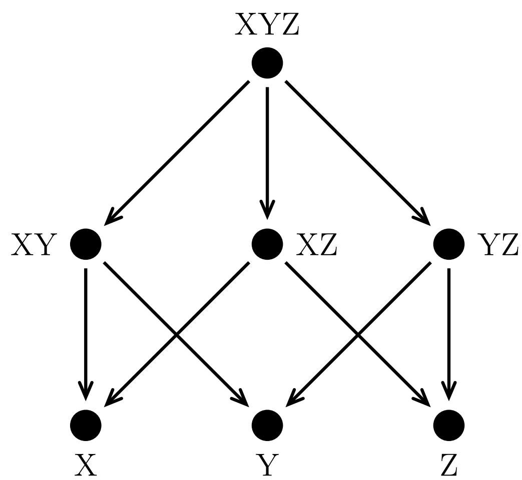

(d) Set-inclusion

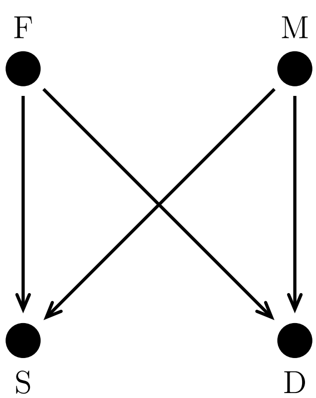

(e) Crown

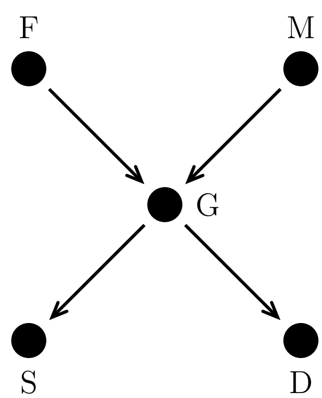

(f) Cross

The hierarchy from Example 5 will play an important role in our analysis. Consequently, the following definition will be useful. Given any hierarchy $H = (N, \geq)$, we say that a subhierarchy $H' = (N', \geq)$ is a *crown subhierarchy* if N' consists of four nodes $\{n_1, n_2, n_3, n_4\}$ and the partial order is defined by $n_1 \geq n_3, n_4$, and $n_2 \geq n_3, n_4$ (with n_1 not comparable with n_2 and n_3 not comparable to n_4).

Example 6 (Cross)**.** The parents from the crown in Example 5 have hired a governess G, who manages the children and reports to the parents. The new partial order is $F, M \geq G$; $G \geq S, D$ and the new *cross* hierarchy is depicted in Figure 1f. Note that the nodes in the crown are a subset of the nodes in the cross we consider here. Moreover, the ordering relation on the crown agrees with the ordering relation on the cross. Yet, the subhierarchy of the cross induced by the subset $\{F, M, S, D\}$ is not a crown subhierarchy since F and M do not cover S and D. As we will see, this observation will play an important role when we consider universal constructibility of the crown and the cross.

3.2 Universal constructibility of the example hierarchies

A chain is universally constructible

Every information hierarchy that is a chain (i.e., that is totally ordered) is universally constructible. To see why, it is helpful to note the following result.[6]

Lemma 1 (Lemma 4 from Gentzkow and Kamenica, 2017)**.** *For any τ, τ', and π s.t. $\tau' \succsim \tau$ and $\langle \pi \rangle = \tau$, $\exists \pi'$ s.t. $\pi' \trianglerighteq \pi$ and $\langle \pi' \rangle = \tau'$.*

In other words, take any signal that induces some distribution of beliefs. There is a refinement of this signal that induces any given more-informative distribution of beliefs.

Now, consider the chain hierarchy H from Example 1. Consider some monotone belief allocation β on H. Given $\beta(D)$, let $\sigma(D)$ be any signal that induces $\beta(D)$. Since $\beta(C) \succsim \beta(D)$, by Lemma 1, there exists some signal $\pi' \trianglerighteq \sigma(D)$ that induces $\beta(C)$; let $\sigma(C) = \pi'$. Similarly, since $\beta(B) \succsim \beta(C)$, there is a $\pi'' \trianglerighteq \sigma(C)$ that induces $\beta(B)$; let $\sigma(B) = \pi''$. Etc.

Since this is the only crown we consider in this paper, we refer to it as *the* crown. Note that "crown" also has distinct meanings in graph theory.

[6]This result first appears as Theorem 1 in Green and Stokey (1978). The statement (and proof) in Gentzkow and Kamenica (2017) uses the same notation as this paper.

A tree is universally constructible

Establishing that a tree is universally constructible is somewhat more complicated than for a chain. The proof relies on the following result, which might be of independent information-theoretic interest:

Lemma 2. *Given π^H, π^L, and τ s.t. $\pi^H \trianglerighteq \pi^L$ and $\langle \pi^L \rangle \succsim \tau$, $\exists \pi^*$ such that (i) $\langle \pi^* \rangle = \tau$, and (ii) $\forall \pi$ s.t. $\pi^H \trianglerighteq \pi \trianglerighteq \pi^L$, $\langle \pi \vee \pi^* \rangle = \langle \pi \rangle$.*

All proofs omitted in the text are in the Appendix. To understand the content of Lemma 2, it is helpful to first note that the following, stronger, conjecture does not hold. One might think that, analogously to Lemma 1, any signal π^L that induces some distribution of beliefs $\langle \pi^L \rangle$ can be coarsened into a signal π^* that induces any given less-informative distribution of beliefs τ: for any τ and π^L s.t. $\langle \pi^L \rangle \succsim \tau$, $\exists \pi^*$ s.t. $\langle \pi^* \rangle = \tau$ and $\pi^L \trianglerighteq \pi^*$. This is not the case.[7] Lemma 2 implies, however, that we can nonetheless always find a π^* s.t. $\langle \pi^* \rangle = \tau$ and $\langle \pi^L \vee \pi^* \rangle = \langle \pi^L \rangle$, even though we cannot guarantee that $\pi^L \trianglerighteq \pi^*$. Moreover, the Lemma further implies we can find a suitable π^* so that $\langle \pi \vee \pi^* \rangle = \langle \pi \rangle$ for any π such that $\pi^L \trianglelefteq \pi \trianglelefteq \pi^H$ (although the choice of π^* depends on π^H).

Now, consider the tree hierarchy H from Example 2. Take some monotone β on H. We construct a monotone σ that induces β as follows. Given $\beta(A1)$, let $\sigma(A1)$ be any signal that induces $\beta(A1)$. We can follow the same procedure as in the case of the chain to (tentatively) assign suitable $\sigma(D1)$ and $\sigma(P)$. Consider assigning a signal to $A2$. The complication is that there may not exist a signal π such that $\langle \pi \rangle = \beta(A2)$ and yet $\sigma(D1) \trianglerighteq \pi$. This is where Lemma 2 comes into play. By Lemma 2, we know there is a signal π^* such that $\langle \pi^* \rangle = \beta(A2)$, $\langle \sigma(D1) \vee \pi^* \rangle = \langle \sigma(D1) \rangle = \beta(D1)$, and $\langle \sigma(P) \vee \pi^* \rangle = \langle \sigma(P) \rangle = \beta(P)$. Thus, we can replace the initial assignment of signals to $D1$ and P with $\sigma(D1) \vee \pi^*$ and $\sigma(P) \vee \pi^*$, respectively. A similar procedure (with repeated reassignment of the previously assigned signals) can then be used to sequentially assign signals to $D2$, $A3$, and $A4$. The details of this procedure, applied to any hierarchy whose graph is a forest, are the heart of the proof of Proposition 1 below.

[7] For example, as pointed out in Gentzkow and Kamenica (2017), suppose that $\langle \pi^L \rangle$ is fully informative, every signal realization in π^L satisfies $p(s|\omega) = 1$ for some ω, and the support of τ has more elements than $|\Omega|$. Then, π must have more elements than π^L and thus we cannot have $\pi^L \trianglerighteq \pi$.

Figure 2: Distributions over $\Omega = \{0,1\}$ showing the diamond hierarchy is not universally constructible

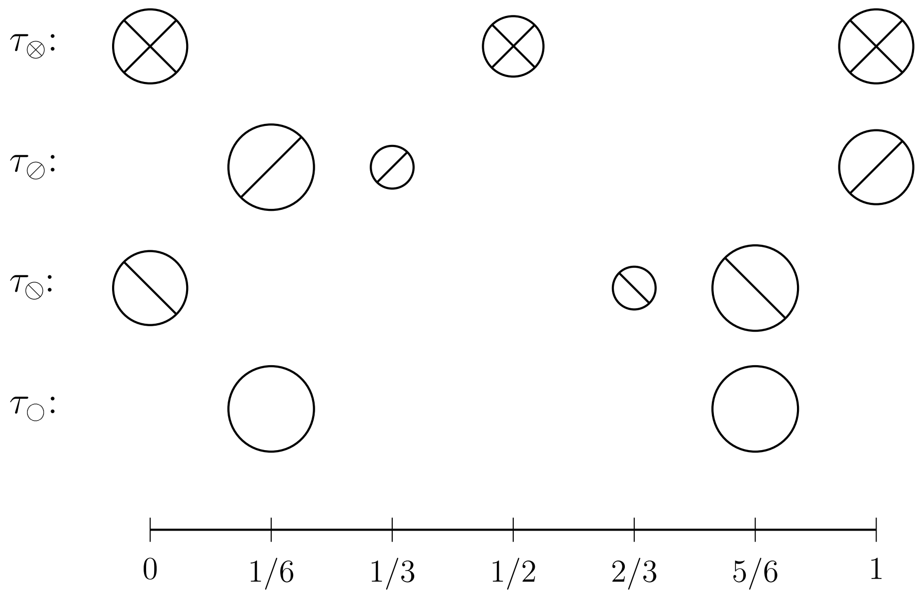

The diamond is not universally constructible

Consider the diamond hierarchy from Example 3. This hierarchy is not 2-universally constructible (and thus is not K-universally constructible for any $K \geq 2$). Suppose $\Omega = \{0,1\}$ with an equiprobable prior. Since the state space is binary, we associate each belief with $\Pr(\omega = 1)$; a belief is depicted as a number in the unit interval.

Consider the belief allocation β that respectively assigns to A, B, C, and D the distributions of beliefs $\tau_{\otimes}$, $\tau_{\oslash}$, $\tau_{\circleddash}$, and $\tau_{\bigcirc}$, as indicated in Figure 2. (We will refer these four distributions of beliefs again below, so it is helpful to give them names.) In the figure, we depict each distribution of beliefs as a collection of circles with matching markings. Each circle represents a belief in the support of the distribution, where the area of a circle is proportional to the probability mass on that belief. Distribution $\tau_{\otimes}$ puts probability $\frac{3}{8}$ on $\mu = 0$, probability $\frac{1}{4}$ on $\mu = \frac{1}{2}$, and probability $\frac{3}{8}$ on $\mu = 1$. Distribution $\tau_{\oslash}$ puts probability $\frac{1}{2}$ on $\mu = \frac{1}{6}$, probability $\frac{1}{8}$ on $\mu = \frac{1}{3}$, and probability $\frac{3}{8}$ on $\mu = 1$. Distribution $\tau_{\circleddash}$ puts probability $\frac{3}{8}$ on $\mu = 0$, probability $\frac{1}{8}$ on $\mu = \frac{2}{3}$, and probability $\frac{1}{2}$ on $\mu = \frac{5}{6}$. Distribution $\tau_{\bigcirc}$ puts probability $\frac{1}{2}$ on $\mu = \frac{1}{6}$ and probability $\frac{1}{2}$ on $\mu = \frac{5}{6}$.

It is easy to see that β is monotone.[8] We will now argue that there does not exist a monotone

[8]To see that $\beta(B) = \tau_{\oslash}$ is a mean-preserving spread of $\beta(D) = \tau_{\bigcirc}$, note that we can obtain $\tau_{\oslash}$ from $\tau_{\bigcirc}$ by spreading the realization $\mu = \frac{5}{6}$ in $\tau_{\bigcirc}$ to $\left\{\frac{1}{3}, 1\right\}$ in $\tau_{\oslash}$ and leaving the realization $\mu = \frac{1}{6}$ in $\tau_{\bigcirc}$ unchanged. To see

signal allocation that induces β.

Consider the joint distribution of beliefs on $\{A, B, C, D\}$ induced by any monotone signal allocation σ. Specifically, consider the conditional probability of $\tilde{\mu}_{\sigma(A)} = 1$ given $\tilde{\mu}_{\sigma(D)} = \frac{1}{6}$. Note that for any two nodes n and n', $\sigma(n) \trianglerighteq \sigma(n')$ implies the martingale property that $\mathbb{E}\left[\tilde{\mu}_{\sigma(n)} | \tilde{\mu}_{\sigma(n')}\right] = \tilde{\mu}_{\sigma(n')}$. Hence, $\sigma(B) \trianglerighteq \sigma(D)$ implies that $\tilde{\mu}_{\sigma(B)} = \frac{1}{6}$ whenever $\tilde{\mu}_{\sigma(D)} = \frac{1}{6}$.[9] Similarly, since $\sigma(A) \trianglerighteq \sigma(B)$, we must have $\tilde{\mu}_{\sigma(A)} = 1$ whenever $\tilde{\mu}_{\sigma(B)} = 1$. Finally, since $Pr\left(\tilde{\mu}_{\sigma(A)} = 1\right) = Pr\left(\tilde{\mu}_{\sigma(B)} = 1\right)$, we must have that $Pr\left(\tilde{\mu}_{\sigma(A)} = 1 | \tilde{\mu}_{\sigma(B)} \neq 1\right) = 0$. Combining these observations, we obtain $Pr\left(\tilde{\mu}_{\sigma(A)} = 1 | \tilde{\mu}_{\sigma(D)} = \frac{1}{6}\right) = 0$.

However, applying a similar logic to the assumption that $\sigma(A) \trianglerighteq \sigma(C) \trianglerighteq \sigma(D)$ tells us that $Pr\left(\tilde{\mu}_{\sigma(C)} = \frac{2}{3} | \tilde{\mu}_{\sigma(D)} = \frac{1}{6}\right) > 0$ and $Pr\left(\tilde{\mu}_{\sigma(A)} = 1 | \tilde{\mu}_{\sigma(C)} = \frac{2}{3} \ \& \ \tilde{\mu}_{\sigma(D)} = \frac{1}{6}\right) > 0$. Combining these two inequalities yields $Pr\left(\tilde{\mu}_{\sigma(A)} = 1 | \tilde{\mu}_{\sigma(D)} = \frac{1}{6}\right) > 0$, which contradicts the conclusion we derived from the fact that $\sigma(A) \trianglerighteq \sigma(B) \trianglerighteq \sigma(D)$. Therefore, no monotone signal allocation can induce β.

Complementary to the preceding discussion, we can give the following informal explanation of why the diamond is not universally constructible. The belief allocation we constructed is such that along each edge of the diamond, there is a unique way to spread the less-informative belief distribution to produce the more-informative belief distribution. In particular, there is a unique conditional distribution over beliefs at B given the realized belief at D, a unique conditional distribution over beliefs at A given a realized belief at B, etc. Thus, given a belief realization at D, we can derive a distribution of belief realizations at A by integrating over beliefs at B or by integrating over beliefs at C. If there were a signal allocation that induced these beliefs, then the conditional distribution of the belief at A given the realized belief at D must be "independent of path" up the diamond. As we showed above, this is not the case. The reason that this is possible is that while the Blackwell ordering requires τ_A to be more informative than τ_B and τ_C in the usual ex ante sense, it does not require the same ranking of *conditional* distributions of beliefs, given the realized belief at D. This latter condition is, however, an implication of the refinement order for beliefs that can be induced by a monotone signal allocation.

that $\beta(A) = \tau_{\otimes}$ is a mean-preserving spread of $\beta(B) = \tau_{\oslash}$, note that we can obtain $\tau_{\otimes}$ from $\tau_{\oslash}$ by spreading the realizations $\mu = \frac{1}{6}$ and $\mu = \frac{1}{3}$ in $\tau_{\oslash}$ to $\left\{0, \frac{1}{2}\right\}$ in $\tau_{\otimes}$ and leaving the realization $\mu = 1$ in $\tau_{\oslash}$ unchanged. The argument for why $\beta(A) \succsim \beta(C) \succsim \beta(D)$ is symmetric.

[9] Since $B \geq D$, monotonicity of σ implies $\sigma(B) \trianglerighteq \sigma(D)$, which in turn guarantees $\mathbb{E}\left[\tilde{\mu}_{\sigma(B)} | \tilde{\mu}_{\sigma(D)}\right] = \tilde{\mu}_{\sigma(D)}$. Since the support of $\beta(B)$ is $\left\{\frac{1}{6}, \frac{1}{3}, 1\right\}$, the only way to have $\mathbb{E}\left[\tilde{\mu}_{\sigma(B)} | \tilde{\mu}_{\sigma(D)} = \frac{1}{6}\right] = \frac{1}{6}$ is to have $\tilde{\mu}_{\sigma(B)} = \frac{1}{6}$ whenever $\tilde{\mu}_{\sigma(D)} = \frac{1}{6}$.

A set-inclusion hierarchy is not universally constructible

Consider the environment, described in Example 4, where a decision maker has access to three unknown sources of information. This environment induced the set-inclusion hierarchy depicted in Figure 1d. Note that the graph of the diamond hierarchy can be seen as a subgraph of the graph in Figure 1d, if we associate A with XYZ, B with XY, C with YZ, and D with Y. Consequently, the set-inclusion hierarchy is not universally constructible for the same reason that the diamond hierarchy is not universally constructible. In particular, any belief allocation β on the set-inclusion hierarchy that assigns $\beta(XYZ) = \tau_{\otimes}$, $\beta(XY) = \tau_{\oslash}$, $\beta(YZ) = \tau_{\circleddash}$, and $\beta(Y) = \tau_{\bigcirc}$ cannot be induced by a monotone signal allocation.[10]

It is worthwhile to note that the mere fact that A, B, C, and D in the diamond are ordered the same way as XYZ, XY, YZ, and Z in the set-inclusion hierarchy does not by itself mean that non-constructibility of the diamond implies non-constructibility of the set-inclusion hierarchy. Specifically, suppose there is some hierarchy $(N, \geq)$ and some hierarchy $(N', \geq')$ such that $N' \subseteq N$, and for any $n, m \in N'$, we have $n \geq m \Leftrightarrow n \geq' m$. It could nonetheless be the case that $(N, \geq)$ is universally constructible even though $(N', \geq')$ is not. The next two examples illustrate this possibility. Section 4.2.1 discusses the issue in detail.

The crown is not universally constructible

Consider the crown hierarchy from Example 5. This hierarchy is not K-universally constructible for $K \geq 3$. (As we discuss below, it is 2-universally constructible). Suppose $\Omega = \{0, 1, 2\}$ with a prior $\mu_0(\omega = 0) = \frac{1}{2}$, $\mu_0(\omega = 1) = \frac{1}{4}$, and $\mu_0(\omega = 2) = \frac{1}{4}$. We represent each belief as a pair (x, y) in the unit square with $x + y \leq 1$, where $Pr(\omega = 1) = x$ and $Pr(\omega = 2) = y$. Consider the belief allocation β that assigns to F, M, S, and D the distributions of belief indicated in Figure 3. As before, we depict each distribution of beliefs as a collection of circles. If a letter $n \in \{F, M, S, D\}$ appears inside a circle, then the belief indicated by this circle is in the support of $\beta(n)$. If a belief is in the support of both $\beta(n)$ and $\beta(n')$, both n and n' appear inside that circle; moreover, these two distributions put the same probability mass on that belief (with the mass indicated by the area

[10]Moreover, it is easy to see that there exists a monotone belief allocation that assigns these distributions of beliefs to these nodes. For example, we could set $\beta(X) = \beta(Z) = \beta(XZ) = \underline{\tau}$.

Figure 3: Distributions over $\Omega = \{0, 1, 2\}$ showing the crown is not universally constructible

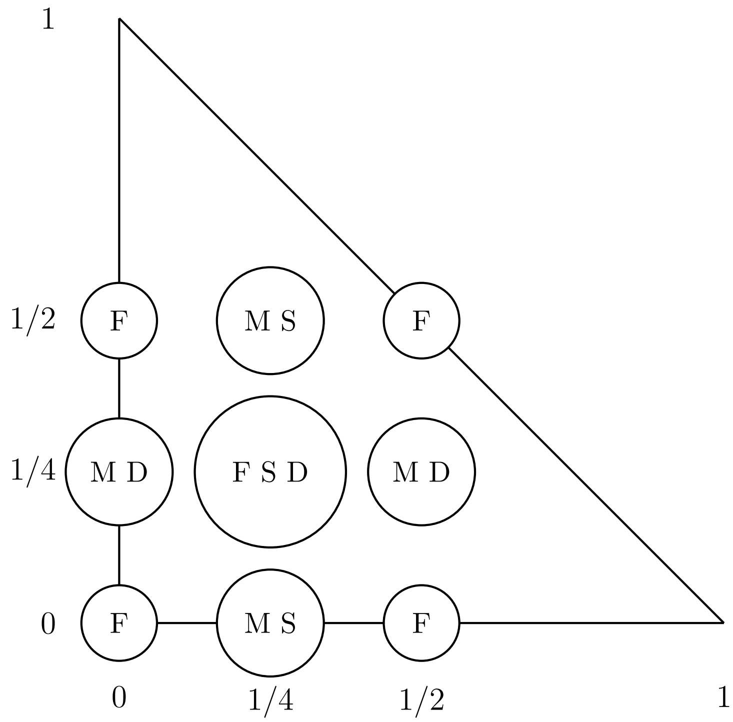

of the circle).

It is easy to see from Figure 3 that β is monotone.[11] We will now argue that there does not exist a monotone signal allocation that induces β. As before, consider the joint distribution of beliefs on $\{F, M, S, D\}$ induced by any monotone signal allocation σ. Specifically, consider the conditional probability of $\tilde{\mu}_{\sigma(S)} = \mu_0$ given $\tilde{\mu}_{\sigma(D)} = \mu_0$. Since $\sigma(M) \trianglerighteq \sigma(S)$, we have that $\tilde{\mu}_{\sigma(S)} = \mu_0 \Leftrightarrow \tilde{\mu}_{\sigma(M)} \in \left\{\left(0, \frac{1}{4}\right), \left(\frac{1}{2}, \frac{1}{4}\right)\right\}$. Since $\sigma(M) \trianglerighteq \sigma(D)$, we have that $\tilde{\mu}_{\sigma(D)} = \mu_0 \Leftrightarrow \tilde{\mu}_{\sigma(M)} \in \left\{\left(\frac{1}{4}, 0\right), \left(\frac{1}{4}, \frac{1}{2}\right)\right\}$. Hence, the joint probability of $\tilde{\mu}_{\sigma(S)} = \mu_0$ and $\tilde{\mu}_{\sigma(D)} = \mu_0$ must be zero, i.e., $Pr\left(\tilde{\mu}_{\sigma(S)} = \mu_0 | \tilde{\mu}_{\sigma(D)} = \mu_0\right) = 0$. But, $\sigma(F) \trianglerighteq \sigma(S)$ implies $\tilde{\mu}_{\sigma(S)} = \mu_0 \Leftrightarrow \tilde{\mu}_{\sigma(F)} = \mu_0$, and $\sigma(F) \trianglerighteq \sigma(D)$ implies $\tilde{\mu}_{\sigma(D)} = \mu_0 \Leftrightarrow \tilde{\mu}_{\sigma(F)} = \mu_0$. Hence, $Pr\left(\tilde{\mu}_{\sigma(S)} = \mu_0 | \tilde{\mu}_{\sigma(D)} = \mu_0\right) = 1$. We have reached a contradiction: no monotone σ can induce the belief allocation β.

For future reference, we summarize this discussion with the following formal result:

[11]To see that $\beta(F)$ is a mean-preserving spread of $\beta(S)$, note that we can obtain $\beta(F)$ from $\beta(S)$ by (i) spreading the realization $\left(\frac{1}{4}, 0\right)$ in $\beta(S)$ to $\left\{(0,0), \left(\frac{1}{2}, 0\right)\right\}$ in $\beta(F)$, (ii) leaving the realization $\left(\frac{1}{4}, \frac{1}{4}\right)$ in $\beta(S)$ unchanged, and (iii) spreading the realization $\left(\frac{1}{4}, \frac{1}{2}\right)$ in $\beta(S)$ to $\left\{\left(0, \frac{1}{2}\right), \left(\frac{1}{2}, \frac{1}{2}\right)\right\}$. To see that $\beta(M)$ is a mean-preserving spread of $\beta(S)$, note that we can obtain $\beta(M)$ from $\beta(S)$ by (i) leaving the realization $\left(\frac{1}{4}, 0\right)$ in $\beta(S)$ unchanged, (ii) spreading the realization $\left(\frac{1}{4}, \frac{1}{4}\right)$ in $\beta(S)$ to $\left\{\left(0, \frac{1}{4}\right), \left(\frac{1}{2}, \frac{1}{4}\right)\right\}$ in $\beta(M)$, and (iii) leaving the realization $\beta(S) = \left(\frac{1}{4}, \frac{1}{2}\right)$ unchanged. The argument for why $\beta(F)$ and $\beta(M)$ are mean-preserving spreads of $\beta(D)$ is symmetric.

Lemma 3. *A crown subhierarchy is not* 3*-universally constructible.*

The cross is universally constructible

The cross hierarchy from Example 6 is universally constructible. The argument is analogous to the argument for why the tree hierarchy from Example 2 is universally constructible. The procedure we discussed in that example for how to construct a monotone σ to induce any monotone β applies to the cross as well.

It is worthwhile to note that, while the nodes in the crown from the previous example are a subset of those in the cross (with the same partial order), the cross is universally constructible even though the crown is not. This might seem puzzling at first. One might think that the impossibility of constructing a monotone σ that induces the belief allocation β on $\{F, M, S, D\}$ considered in the previous example means that the cross is also not universally constructible. But, even though β is monotone with respect to the crown, there is no way of extending that belief allocation to the cross (by assigning some $\beta(G)$ to node G) in a way that would preserve monotonicity.

That the cross is universally constructible does, however, imply that the crown is 2-universally constructible. This follows from the fact that the set of distributions of beliefs under the Blackwell order is a lattice when the state space is binary (Kertz and Rösler, 2000; Müller and Scarsini, 2006). Consider any monotone belief allocation β on $\{F, M, S, D\}$. If the state space is binary, the lattice property implies there exists a unique distribution of beliefs $\beta(S) \vee \beta(D)$ such that $\beta(S) \vee \beta(D) \succsim \beta(S), \beta(D)$ and $\tau \succsim \beta(S) \vee \beta(D)$ for any $\tau \succsim \beta(S), \beta(D)$. Now, let $\hat{\beta}$ be the belief allocation on the cross that sets $\hat{\beta}(n) = \beta(n)$ for $n \in \{F, M, S, D\}$ and $\hat{\beta}(G) = \beta(S) \vee \beta(D)$. It is immediate that $\hat{\beta}$ is monotone. By the universal constructibility of the cross, there is a monotone $\hat{\sigma}$ that induces $\hat{\beta}$. Restricting $\hat{\sigma}$ to $\{F, M, S, D\}$ yields a monotone signal allocation on the crown that induces β. Hence, the crown is 2-universally constructible.

4 Universal constructibility

4.1 Main result

We now present our main result, a characterization of universally constructible hierarchies.

Theorem 1. *An information hierarchy is universally constructible if and only if its undirected graph is a forest.*

The proof of Theorem 1 is broken up into two propositions, which separately establish sufficiency and necessity of the forest condition.

Proposition 1. *An information hierarchy is universally constructible if its undirected graph is a forest.*

Since a forest is a union of disjoint trees, this result is a straightforward consequence of the aforementioned fact that trees are universally constructible. We have already given some intuition for the universal constructibility of trees in Section 3. A formal proof of Proposition 1 is in the Appendix.

Proposition 2. *Given $K \geq 3$, an information hierarchy is K-universally constructible only if its undirected graph is a forest.*

A rigorous proof of Proposition 2 is in the Appendix. The argument is lengthy and introduces several new concepts. The remainder of this section provides an outline of the proof.

4.2 Outline of the proof of Proposition 2

Recall that in Section 3, we presented two examples of hierarchies which are not universally constructible, the diamond and the crown. Specifically, we constructed monotone belief allocations which are not induced by any monotone signal allocation. These constructions turn out to be "canonical" in the sense that for any hierarchy that is not universally constructible, we can generalize these belief allocations to establish non-universal constructibility.

4.2.1 Constructibility and closed subhierarchies

Sometimes, we can establish that a hierarchy is not universally constructible by noting that it has a subhierarchy that is not universally constructible. As a simple example, suppose we take the diamond in Figure 1c, and form a new hierarchy H' by adding nodes that are above A or below D. Then, we can extend any monotone belief allocation β on the diamond to H' by assigning full information to the nodes above A and no information to the nodes below B. The resulting belief allocation β' will be monotone on H', and it is constructible if and only if β was constructible on the diamond. Thus, any hierarchy which "embeds" the diamond in this sense is not K-constructible for any $K \geq 2$.

A natural conjecture might therefore be that a hierarchy is not universally constructible if it contains a subhierarchy that is not universally constructible. Without further conditions, this conjecture is false. This was previously demonstrated by the crown and the cross: the cross is universally constructible, but if we drop its center node, the resulting subhierarchy is a crown subhierarchy, which is not universally constructible. The conjecture is true, however, if we add an additional hypothesis on the subhierarchy. Given $H = (N, \geq)$, a subhierarchy $H' = (N', \geq)$ is *closed* if for every $n', n'' \in N'$ and $n \in N$, $n' \geq n \geq n''$ implies $n \in N'$. In other words, N' contains all the nodes from N that are between the nodes of N'.

Lemma 4 in the Appendix establishes that a hierarchy is universally constructible only if its closed subhierarchies are universally constructible. The basic idea behind the proof is as follows. Consider some hierarchy H that is universally constructible. Let H' be a closed subhierarchy and let β' be a monotone belief allocation on H'. We will extend β' to a monotone belief allocation β on H. Since H' is closed, any node in H that is not in H' falls into one of three mutually exclusive categories: (i) it is above some node in H', (ii) it is below some node in H', or (iii) it is not comparable with any node in H'. In case (i), we allocate full information to the node, and in cases (ii) and (iii), we allocate no information to the node. The resulting β is monotone. Moreover, since H is universally constructible, there is a monotone signal allocation σ that induces β. The restriction of σ to H' is monotone and induces β'. Since β' was arbitrary, we conclude that H' is universally constructible.

4.2.2 Unions of non-comparable paths

We now introduce an important class of hierarchies, termed unions of non-comparable paths (UNPs). We adapt the belief construction for the diamond to establish that UNPs are not universally constructible. UNPs do not necessarily embed the diamond, so the proof method does not rely on the observation from the previous subsection.

Figure 4: An example of a UNP. Here, the curves depict undirected paths between A' and D', and slopes of the curve denote "local" comparisons between nodes. For example, the path through B' is decreasing, indicating that it is a directed path with B' below A' and above D'. In contrast, the paths through C' and C'' are not directed; while C' is above D', it is non-comparable with A'. This a UNP because it satisfies three properties: (i) node A' is maximal; (ii) node D' is minimal; and (iii) a pair of nodes are comparable only if they are in the same undirected path.

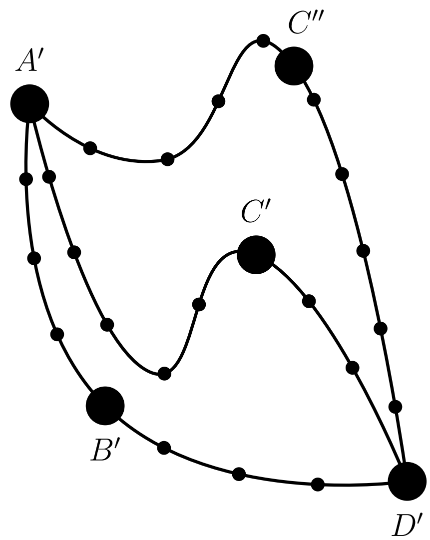

Consider a hierarchy H whose graph is of the form depicted in Figure 4. In particular, the graph is a union of (at least two) undirected paths between the pair of nodes A' and D' satisfying three properties: (i) node A' is maximal; (ii) node D' is minimal; and (iii) a pair of nodes are comparable only if they are in the same undirected path. Note that we have not made any assumptions about the ordering within paths. We refer to such a hierarchy as a *union of non-comparable paths (UNP).*

Lemma 8 in the Appendix shows that UNPs are not 2-universally constructible. The argument is closely linked to the argument for non-universal constructibility of the diamond.

We first establish that if a belief allocation on a UNP allocates the same belief distribution on any pair of nodes n and n' that are in the same path (neither node being A' or D'), then any monotone signal allocation that induces it yields beliefs at n and n' that are perfectly correlated.

Now, consider some belief allocation β on a UNP H based on the belief allocation on the diamond from Section 3. In particular, set $\beta(A') = \tau_{\otimes}$; $\beta(D') = \tau_{\bigcirc}$; for nodes n in one of the paths of H, say the path including B' in Figure 4 (excluding A' and D'), set $\beta(n) = \tau_{\oslash}$; and for all remaining nodes n (namely the nodes in the paths including C' and C'' in Figure 4) set $\beta(n) = \tau_{\obslash}$.

We can then show that β cannot be induced by a monotone signal allocation. The argument is somewhat subtle and relies on the particular belief allocation we consider. Roughly speaking, the nodes in the path with B' (whose belief realizations are perfectly correlated) collectively serve the role of node B in the diamond, and the other paths serve the role of node C from the diamond.[12]

4.2.3 Minimal cyclic closed (MCC) subhierarchies

The observations from the previous two subsections imply that to complete the proof of Proposition 2, it suffices to establish the result that if a hierarchy is not a forest, then it contains a closed subhierarchy that is either a UNP or a crown subhierarchy. We show this result as follows.

Suppose some hierarchy H is not a forest. Then H contains a cycle; moreover, since any hierarchy is itself closed, H contains a subhierarchy that is cyclic and closed. It is then straightforward to argue that H must also contain a *minimal cyclic closed subhierarchy (MCC)*, i.e., a subhierarchy $H' = (N', \geq)$ such that: (i) H' is cyclic and closed, and (ii) there exists no subhierarchy $H'' = (N'', \geq)$ with $N'' \subsetneq N'$ such that H'' is cyclic and closed.

The second step is to show that any MCC must be either a UNP or a crown subhierarchy. This result takes considerable effort to prove formally, but the basic idea is as follows. An MCC H' contains some maximal nodes, which are not covered by any other node, and minimal nodes, which do not cover any other nodes.

One possibility is that every maximal node in H' covers every minimal node. In this case, the minimality condition, together with the existence of a cycle, implies that the subhierarchy must contain exactly four nodes, and in fact is a crown subhierarchy.

Alternatively, there is a maximal node $\overline{n}$ and a minimal node $\underline{n}$ that do not cover each other. Then, H' must be a UNP. To see this, note that there are two cases. One case is that H' is simply the set of nodes that are between $\overline{n}$ and $\underline{n}$, i.e., a *between set*. Then, H' consists of a series of directed paths between $\overline{n}$ and $\underline{n}$. Now, if nodes in distinct paths were comparable, then it would be possible to find a smaller cyclic closed subhierarchy, as illustrated in Figure 5. That would violate

[12]There exist belief allocations on the diamond whose "extension" to H is nonetheless constructible. (An example is available from the authors upon request.) The reason why this is possible is that we have $A \geq D$, but we do not require that $A' \geq D'$; thus signal allocations $\hat{\sigma}$ on the diamond and σ on H with $\hat{\sigma}(A) = \sigma(A')$ and $\hat{\sigma}(D) = \sigma(D')$ might be monotone on H but not on the diamond. What is special about the belief allocation we considered on the diamond is that any signal allocation σ that induces an extension of it on H must satisfy $\sigma(A') \trianglerighteq \sigma(D')$.

Figure 5: A between set which is not an MCC. Since the nodes A and B are comparable, there is a smaller closed cyclic subhierarchy, namely the nodes that are between A and $\underline{n}$.

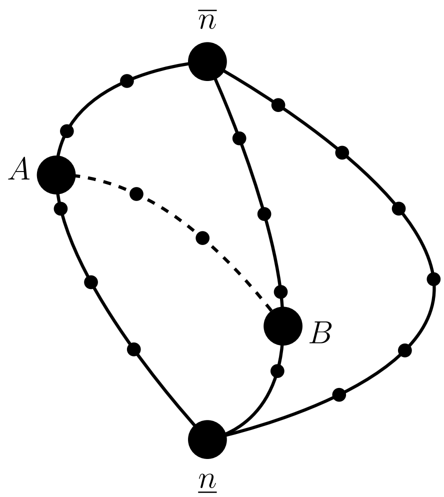

minimality. Thus, nodes must not be comparable across paths and H' is a UNP. The second case is that H' is not a between set. Then, every cycle in H' must contain every node in N'. (This is a non-trivial, technical result that is established in Lemma 5 in the Appendix.) Any such "spanning cycle" can be decomposed into two undirected paths between a maximal node and a minimal node. If any nodes in these two paths were comparable, we could find a smaller cycle that does not contain every node in N'. Since this cannot happen, H' is a UNP.

With this result, we have completed the outline of the proof of Proposition 2.

5 Discussion

5.1 Rationalizing reaction to unknown sources of information

Consider some agent who obtains information from multiple sources. If we do not know the information-generating process, what restrictions does the agent's rationality impose on her potential reactions to this information? Concretely, suppose Anne is a decision maker with access to a set of Blackwell experiments $\{x_1, x_2, ..., x_M\}$. Suppose we see the distribution of Anne's beliefs that arises after she observes any subset of these experiments. Specifically, our dataset $\mathcal{D} = \{\tau_S\}_{S \subseteq \{x_1,...,x_M\}}$ tells us the distribution of Anne's beliefs for every non-empty subset of experiments that she observes. When can we rationalize a given dataset $\mathcal{D}$ in the sense that we can associate each experiment x_i with some signal (i.e., an element of Π) and conclude that Anne's belief formation is consistent with Bayes' rule?

To be rationalized, belief distributions in $\mathcal{D}$ have to satisfy two obvious properties. First, there is *Bayes plausibility*: the average belief cannot differ across sets of experiments, i.e., $\mathbb{E}_{\tau_S}[\mu] = \mathbb{E}_{\tau_{S'}}[\mu]$ for any two subsets S and S'. Second, there is *Blackwell monotonicity*: observing a larger set of experiments necessarily induces a more dispersed distribution of beliefs, i.e., τ_S is a mean-preserving spread of $\tau_{S'}$ if $S' \subseteq S$. A natural question is whether these are the only properties imposed by Bayesian updating.

Theorem 1 tells us that the answer is No. When there are three or more experiments,[13] Bayesian updating requires more than just Bayes plausibility and Blackwell monotonicity. To see why, consider the set-inclusion information hierarchy H where each non-empty collection of experiments $S \subseteq \{x_1, ..., x_M\}$ is associated with a node n_S and the partial order is the superset order: $n_S \geq n_{S'}$ if $S' \subseteq S$. As illustrated in Figure 1d, the undirected graph of this information hierarchy H contains a cycle. By Theorem 1, this means that there is some monotone belief allocation on H, say β, that cannot be induced by any monotone signal allocation on H. Now, we can associate with this β a dataset $\mathcal{D} = \{\tau_S\}_{S \subseteq \{x_1,...,x_M\}}$ by setting $\tau_S = \beta(n_S)$ and note that $\mathcal{D}$ necessarily satisfies Bayes plausibility and Blackwell monotonicity (since β is monotone). If we could rationalize $\mathcal{D}$ by associating each x_i with some signal $\pi(x_i) \in \Pi$, then the signal allocation $\sigma(n_s) = \bigvee_{x_i \in S} \pi(x_i)$ would induce β and yet be monotone (since $S' \subseteq S$ implies $\bigvee_{x_i \in S} \pi(x_i) \trianglerighteq \bigvee_{x_i \in S'} \pi(x_i)$). This would contradict Theorem 1. Thus, we know that there are datasets that satisfy Bayes plausibility and Blackwell monotonicity, yet cannot be rationalized.

A potentially fruitful direction for future research would be to fully characterize which reactions to unknown sources of information are rationalizable.

5.2 Information design in organizations

One important aspect of designing organizations is deciding how much information to provide (about individuals' prospects for promotion, about the overall performance of the organization,

[13] When there are only two experiments, $M = 2$, it is easy to show that the answer is indeed affirmative. Any reaction to two unknown sources of information that satisfies Bayes plausibility and Blackwell monotonicity is consistent with Bayesian updating.

etc.) to each member of the organization. It is often suboptimal to provide full transparency and share full information with everyone (Fuchs, 2007; Jehiel, 2015; Smolin, 2017). A natural constraint that an information designer might face is that anyone in the organization ought to have (access to) all of the information that is available to her subordinates. Our results tell us that the nature of this constraint interacts with the organization structure.

An organization is said to be hierarchical if every individual, except one (the head of the organization), is subordinate to a single other individual. Proposition 1 implies that, if an organization is hierarchical, the aforementioned constraint can always be satisfied as long as individuals who are higher up in the organization are more informed in the Blackwell sense. In case of some other organizational forms,[14] however, there could be desirable allocations of information which are incompatible with the constraint, even though they provide (Blackwell) more information to those higher up in the organization.

6 Conclusion

The purpose of this paper has been to study the relationship between the Blackwell and the refinement order in a general-purpose model of distributed information. Take some information hierarchy, i.e., a specification of which elements should be more informed. We analyze whether every belief allocation on this hierarchy that is monotone in the Blackwell order (higher elements know more about the state of the world) is compatible with a signal allocation that is monotone in the refinement order (higher elements know everything lower elements know). Our main result is that the answer is affirmative precisely if the undirected graph of the information hierarchy is a forest.

Our work suggests at least two interesting topics for future work. First, our result on the necessity of the forest condition for universal constructibility requires at least three states. A natural question is: which information hierarchies are 2-universally constructible? The answer must be non-trivial, since the crown is 2-universally constructible, while the diamond is not. Second, our analysis has focused on whether *every* monotone belief allocation can be induced by a monotone signal allocation. A natural goal would be to characterize, given an arbitrary information hierarchy, the set of monotone belief allocations that can be induced by a monotone signal allocation.

[14]For instance, suppose that the CEO oversees two middle managers who share the oversight of an employee.

References

Bergemann, D. and Morris, S. (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, **11** (2), 487–522.

Blackwell, D. (1951). Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, The Regents of the University of California.

— (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, pp. 265–272.

Börgers, T., Hernando-Veciana, A. and Krähmer, D. (2013). When are signals complements or substitutes? *Journal of Economic Theory*, **148** (1), 165–195.

Curello, G. and Sinander, L. (2019). *The Lattice Structure of Preference Comparison*. Tech. rep., University of Oxford and Northwestern University.

Fuchs, W. (2007). Contracting with repeated moral hazard and private evaluations. *American Economic Review*, **97** (4), 1432–1448.

Gentzkow, M. and Kamenica, E. (2017). Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, **104**, 411–429.

Green, J. R. and Stokey, N. L. (1978). Two representations of information structures and their comparisons.

Jehiel, P. (2015). Investment strategy and selection bias: An equilibrium perspective on overconfidence.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, **101** (6), 2590–2615.

Kertz, R. P. and Rösler, U. (2000). Complete lattices of probability measures with applications to martingale theory. *Lecture Notes-Monograph Series*, pp. 153–177.

Lu, L. (2014). On crown-free families of subsets. *Journal of Combinatorial Theory, Series A*, **126**, 216–231.

Müller, A. and Scarsini, M. (2006). Stochastic order relations and lattices of probability measures. *SIAM Journal on Optimization*, **16** (4), 1024–1043.

Smolin, A. (2017). *Dynamic Evaluation Design*. Tech. rep., University Library of Munich, Germany.

A Omitted Proofs

A.1 Proof of Lemma 2

Proof of Lemma 2. Let $\hat{\pi}$ be a signal s.t. $\langle\hat{\pi}\rangle = \tau$. Since $\langle\pi^L\rangle \succsim \tau$, there exists a garbling $g : \pi^L \times \hat{\pi} \to [0,1]$ such that $\sum_{\hat{s}\in\hat{\pi}} g(s,\hat{s}) = 1$ $\forall s \in \pi^L$, and $p(\hat{s}|\omega) = \sum_{s\in\pi^L} g(s,\hat{s})\, p(s|\omega)$. For every $\overline{s} \in \pi^H$, let $\underline{s}(\overline{s})$ denote the element of π^L s.t. $\overline{s} \subseteq \underline{s}(\overline{s})$. (This element exists since $\pi^H \trianglerighteq \pi^L$.) Now, $\forall \overline{s} \in \pi^H$, let $\left\{X^{\overline{s}}_{\hat{s}}\right\}_{\hat{s}\in M^{\overline{s}}}$ be a partition of $\overline{s}$ s.t. $\forall\omega$, $\lambda\left(\left\{x|\,(x,\omega) \in X^{\overline{s}}_{\hat{s}}\right\}\right) = \lambda\left(\{x|\,(x,\omega)\in\overline{s}\}\right) g\left(\underline{s}(\overline{s}),\hat{s}\right)$, where $M^{\overline{s}} = \{\hat{s}|g\left(\underline{s}(\overline{s}),\hat{s}\right) > 0\}$. Such a partition exists because $\sum_{\hat{s}\in\hat{\pi}} g\left(\underline{s}(\overline{s}),\hat{s}\right) = 1$ for all $\underline{s}(\overline{s}) \in \pi^L$. Let $\pi^* = \left\{Z^{\hat{s}}\right\}_{\hat{s}\in\hat{\pi}}$ with $Z^{\hat{s}} = \bigcup_{\overline{s}\in\pi^H \text{ s.t.} \hat{s}\in M^{\overline{s}}} X^{\overline{s}}_{\hat{s}}$. We now show that π^* satisfies (i) and (ii). To show (i), it suffices to show that $p\left(Z^{\hat{s}}|\omega\right) = p(\hat{s}|\omega)$ for every $\hat{s}$ and ω. We have

$$
\begin{aligned}
p\left(Z^{\hat{s}}|\omega\right) &= \lambda\left(\left\{x|\,(\omega,x) \in \bigcup_{\overline{s}\in\pi^H \text{ s.t.} \hat{s}\in M^{\overline{s}}} X^{\overline{s}}_{\hat{s}}\right\}\right) \\
&= \sum_{\overline{s}\in\pi^H \text{ s.t.} \hat{s}\in M^{\overline{s}}} \lambda\left(\left\{x|\,(x,\omega) \in X^{\overline{s}}_{\hat{s}}\right\}\right) \\
&= \sum_{\overline{s}\in\pi^H \text{ s.t.} \hat{s}\in M^{\overline{s}}} \lambda\left(\{x|\,(x,\omega)\in\overline{s}\}\right) g\left(\underline{s}(\overline{s}),\hat{s}\right) \\
&= \sum_{\overline{s}\in\pi^H} \lambda\left(\{x|\,(x,\omega)\in\overline{s}\}\right) g\left(\underline{s}(\overline{s}),\hat{s}\right) \\
&= \sum_{s\in\pi^L} \sum_{\overline{s} \text{ s.t. } \underline{s}(\overline{s})=s} \lambda\left(\{x|\,(x,\omega)\in\overline{s}\}\right) g\left(\underline{s}(\overline{s}),\hat{s}\right) \\
&= \sum_{s\in\pi^L} g(s,\hat{s}) \sum_{\overline{s} \text{ s.t. } \underline{s}(\overline{s})=s} \lambda\left(\{x|\,(x,\omega)\in\overline{s}\}\right) \\
&= \sum_{s\in\pi^L} g(s,\hat{s})\, \lambda\left(\{x|\,(x,\omega)\in s\}\right) \\
&= \sum_{s\in\pi^L} g(s,\hat{s})\, p(s|\omega) \\
&= p(\hat{s}|\omega)\,.
\end{aligned}
$$

To show (ii), consider some π s.t. $\pi^H \trianglerighteq \pi \trianglerighteq \pi^L$ and some $s \in \pi$. Since $\pi \vee \pi^* \trianglerighteq \pi$, there is a partition of s, say $\{s^{\vee}_i\}_{i\in I}$ s.t. $s^{\vee}_i \in \pi \vee \pi^*$ for all i. It will suffice to show that for every ω, ω', and $s^{\vee}_i$, we have

$$\frac{p\left(s^{\vee}_i|\omega\right)}{p\left(s^{\vee}_i|\omega'\right)} = \frac{p(s|\omega)}{p(s|\omega')}.$$

Consider some $s_i^{\vee}$. Note that there exists $\underline{s} \in \pi^L$ with $s \subseteq \underline{s}$ since $\pi \trianglerighteq \pi^L$. Let $Q = \left\{ \overline{s} \in \pi^H | \overline{s} \subseteq s \right\}$. Since $\pi^H \trianglerighteq \pi$, for every ω, $\lambda\left(x | \left(x, \omega\right) \in s\right) = \sum_{\overline{s} \in Q} \lambda\left(x | \left(x, \omega\right) \in \overline{s}\right)$. Note that $\overline{s} \subseteq \underline{s}$ for all $\overline{s} \in Q$. Now, we know that $s_i^{\vee} = s^* \cap s$ for some $s^* \in \pi^*$. By definition of π^*, we know that $s_i^* = \underset{\overline{s} \in \pi^H}{} \underset{\text{s.t.} \hat{s} \in M^{\overline{s}}}{\cup} X_{\hat{s}}^{\overline{s}}$ for some $\hat{s} \in \hat{\pi}$. Hence,

$$
\begin{aligned}
s_i^{\vee} &= \left(\underset{\overline{s} \in \pi^H}{} \underset{\text{s.t.} \hat{s} \in M^{\overline{s}}}{\cup} X_{\hat{s}}^{\overline{s}} \right) \cap s \\
&= \underset{\overline{s} \in \pi^H}{} \underset{\text{s.t.} \hat{s} \in M^{\overline{s}}}{\cup} \left(X_{\hat{s}}^{\overline{s}} \cap s \right) \\
&= \underset{\overline{s} \in Q}{} \underset{\text{s.t.} \hat{s} \in M^{\overline{s}}}{\cup} X_{\hat{s}}^{\overline{s}},
\end{aligned}
$$

where the last equality follows from the fact that $X_{\hat{s}}^{\overline{s}} \subseteq \overline{s}$, and hence $X_{\hat{s}}^{\overline{s}} \cap s = X_{\hat{s}}^{\overline{s}}$ if $\overline{s} \in Q$ and $X_{\hat{s}}^{\overline{s}} \cap s$ is empty if $\overline{s} \notin Q$. Hence,

$$
\begin{aligned}
p\left(s_i^{\vee} | \omega\right) &= \lambda\left(\left\{x | \left(x, \omega\right) \in s_i^{\vee}\right\}\right) \\
&= \lambda\left(\left\{x | \left(x, \omega\right) \in \underset{\overline{s} \in Q}{} \underset{\text{s.t.} \hat{s} \in M^{\overline{s}}}{\cup} X_{\hat{s}}^{\overline{s}}\right\}\right) \\
&= \sum_{\overline{s} \in Q \text{ s.t.} \hat{s} \in M^{\overline{s}}} \lambda\left(\left\{x | \left(x, \omega\right) \in X_{\hat{s}}^{\overline{s}}\right\}\right) \\
&= \sum_{\overline{s} \in Q \text{ s.t.} \hat{s} \in M^{\overline{s}}} \lambda\left(\left\{x | \left(x, \omega\right) \in \overline{s}\right\}\right) g\left(\underline{s}, \hat{s}\right) \\
&= \sum_{\overline{s} \in Q} \lambda\left(\left\{x | \left(x, \omega\right) \in \overline{s}\right\}\right) g\left(\underline{s}, \hat{s}\right) \\
&= g\left(\underline{s}, \hat{s}\right) \sum_{\overline{s} \in Q} \lambda\left(\left\{x | \left(x, \omega\right) \in \overline{s}\right\}\right) \\
&= g\left(\underline{s}, \hat{s}\right) \lambda\left(x | \left(x, \omega\right) \in s\right) \\
&= g\left(\underline{s}, \hat{s}\right) p\left(s | \omega\right).
\end{aligned}
$$

Hence,

$$
\frac{p\left(s_i^{\vee} | \omega\right)}{p\left(s_i^{\vee} | \omega'\right)} = \frac{g\left(\underline{s}, \hat{s}\right) p\left(s | \omega\right)}{g\left(\underline{s}, \hat{s}\right) p\left(s | \omega'\right)} = \frac{p\left(s | \omega\right)}{p\left(s | \omega'\right)},
$$

which completes the proof of Lemma 2. □

A.2 Proof of Proposition 1

Proof of Proposition 1. Suppose $\tilde{G}(H)$ is a forest. Let β be a monotone belief allocation on H. We will construct a monotone signal allocation that induces β. We will construct the hierarchy node by node; each time we add a node, we will assign a signal to the node that is added, and sometimes reassign the signal attached to previously added nodes.

Let us start with some notation and terminology. A *construction procedure* f for hierarchy H is a bijection from $\{1, ..., |N|\}$ to N that specifies the order in which the nodes are added during the construction. Denote $\{f(1), ..., f(l)\}$, i.e., the first l elements added, by N_l^f. If $f(l) = n$, we say that n was added at time l, and we refer to N_{l-1}^f as the previously added nodes. For any subset $N' \subseteq N$, let $CoveredBy(N') = \{n \in N \setminus N' | \exists n' \in N' \text{ that covers } n\}$, $Covering(N') = \{n \in N \setminus N' | \exists n' \in N' \text{ that is covered by } n\}$, and

$$Disconnected\left(N'\right) = \left\{n \in N \setminus N' | \nexists n' \in N' \text{ s.t. there is a path from } n \text{ to } n' \text{ in } G(\tilde{H})\right\}.$$

Now, consider a construction procedure f of the following form. Let $f(1)$ be any element of N. For $l \in \{2, 3, ..., |N|\}$, we let $f(l)$ be an arbitrary element of $CoveredBy\left(N_{l-1}^f\right) \cup Covering\left(N_{l-1}^f\right) \cup Disconnected\left(N_{l-1}^f\right)$. Note that for any $N' \subsetneq N$, $CoveredBy(N') \cup Covering(N') \cup Disconnected(N')$ is not empty.

Claim 1. For each $l \geq 2$, there is at most one edge in $\tilde{G}(H)$ between $f(l)$ and nodes in N_{l-1}^f.

Proof of Claim 1. Suppose toward contradiction that $f(l)$ has an edge in $\tilde{G}(H)$ with distinct $n,n' \in N_{l-1}^f$. Since n and n' connect to $f(l)$, they must belong to the same tree in $\tilde{G}(H)$. Moreover, there must be a path between n and n' in $\tilde{G}\left(\left(N_{l-1}^f, \geq\right)\right)$. To see this, let $\underline{n}$ be the node that was added earliest to N_{l-1}^f among the nodes in the tree to which n and n' belong. For every other node $f(l') \in N_{l-1}^f$ from this tree, we must have $f(l') \in CoveredBy\left(N_{l'-1}^f\right) \cup Covering\left(N_{l'-1}^f\right)$, which in turn means that there is a path from $f(l')$ to $\underline{n}$ in $\tilde{G}\left(\left(N_{l'-1}^f, \geq\right)\right)$ and thus in $\tilde{G}\left(\left(N_{l-1}^f, \geq\right)\right)$. Hence, there is a path from both n and n' to $\underline{n}$ and thus a path between n and n' in $\tilde{G}\left(\left(N_{l-1}^f, \geq\right)\right)$. So, there must be a path between n and n' in $\tilde{G}(H)$ that does not go through $f(l)$. But, because $f(l)$ has an edge with both n and n', there is another path from n to n' that goes through $f(l)$. However, $\tilde{G}(H)$ is a forest, so there cannot be multiple paths between two nodes; we have reached a contradiction. $\diamondsuit$

Now, given this construction procedure f, we assign signals to nodes as follows. At step l, we expand N_{l-1}^f to $N_l^f = N_{l-1}^f \cup f(l)$ and assign signals according to $\sigma^l : N_l^f \to \Pi$. We proceed by induction and show that, as long as the signals previously allocated to nodes in N_{l-1}^f induce appropriate beliefs (i.e., for all $m \in N_{l-1}^f$, $\langle \sigma^{l-1}(m) \rangle = \beta(m)$) and are monotone (i.e., for any $m, m' \in N_{l-1}^f$ such that $m \geq m'$, we have $\sigma^{l-1}(m) \trianglerighteq \sigma^{l-1}(m')$), the σ^l we specify induces appropriate beliefs and is monotone on N_l^f.

First, to node $f(1)$, we assign an arbitrary signal $\sigma^1(f(1))$ such that $\langle \sigma^1(f(1)) \rangle = \beta(f(1))$. Note we are vacuously satisfying the base case of the induction argument: the signal allocation to the single node in N_1^f induces appropriate beliefs and is monotone. For $l \geq 2$, there are three cases: $f(l) \in \mathit{CoveredBy}\left(N_{l-1}^f\right)$, $f(l) \in \mathit{Covering}\left(N_{l-1}^f\right)$, and $f(l) \in \mathit{Disconnected}\left(N_{l-1}^f\right)$.

We first consider the case $f(l) \in \mathit{Covering}\left(N_{l-1}^f\right)$. Note that, by Claim 1, $f(l)$ covers exactly one node in N_{l-1}^f (call this node $\bar{m}$) and is not covered by any nodes in N_{l-1}^f. Since $\beta(f(l)) \succsim \beta(\bar{m})$, there exists some $\pi \trianglerighteq \sigma^{l-1}(\bar{m})$ such that $\langle \pi \rangle = \beta(f(l))$ (cf: Lemma 1). We set $\sigma^l(f(l)) = \pi$ and we keep the signal allocation to nodes in N_{l-1}^f unchanged, i.e., $\sigma^l(m) = \sigma^{l-1}(m)$ for all $m \in N_{l-1}^f$. It is clear that σ^l induces appropriate beliefs (by the inductive hypothesis for $m \in N_{l-1}^f$ and by construction for $f(l)$). We also need to show that this signal allocation on N_l^f is monotone. Consider any $m, m' \in N_l^f$ such that $m > m'$. Since $f(l) \in \mathit{Covering}\left(N_{l-1}^f\right)$, either $m, m' \in N_{l-1}^f$ or $f(l) = m$. In the former case, we know $\sigma^l(m) = \sigma^{l-1}(m) \trianglerighteq \sigma^{l-1}(m') = \sigma^l(m')$ by the inductive hypothesis. If $f(l) = m$, we know $f(l) > \bar{m} \geq m'$. By the inductive hypothesis, $\sigma^l(\bar{m}) = \sigma^{l-1} \trianglerighteq \sigma^{l-1}(m') = \sigma^l(m')$ and thus $\sigma^l(f(l)) \trianglerighteq \sigma^l(\bar{m}) \trianglerighteq \sigma^l(m')$. That completes the proof for this case.

Now consider the case where $f(l) \in \mathit{CoveredBy}\left(N_{l-1}^f\right)$. Let $\underline{m}$ be the node in N_{l-1}^f that covers $f(l)$. Denote $\tau = \beta(f(l))$, $\pi^L = \sigma^{l-1}(\underline{m})$, and $\pi^H = \bigvee_{m \in N_{l-1}^f} \sigma^{l-1}(m)$. By Lemma 2, we know $\exists \pi^*$ such that (i) $\langle \pi^* \rangle = \tau$, and (ii) $\forall \pi$ s.t. $\pi^H \trianglerighteq \pi \trianglerighteq \pi^L$, $\langle \pi \vee \pi^* \rangle = \langle \pi \rangle$. We set $\sigma^l(f(l)) = \pi^*$. For $m \in N_{l-1}^f$, if $m \geq f(l)$, we set $\sigma^l(m) = \sigma^{l-1}(m) \vee \pi^*$; otherwise, we set $\sigma^l(m) = \sigma^{l-1}(m)$. We need to show that σ^l induces appropriate beliefs and is monotone. We have that $\langle \sigma^l(f(l)) \rangle = \langle \pi^* \rangle = \tau = \beta(f(l))$. For $m \in N_{l-1}^f$, first consider cases where $m \geq f(l)$, so $\langle \sigma^l(m) \rangle = \langle \sigma^{l-1}(m) \vee \pi^* \rangle$. Since $m \geq \underline{m}$ (recall that $\underline{m}$ covers $f(l)$), by the inductive hypothesis, $\sigma^{l-1}(m) \trianglerighteq \sigma^{l-1}(\underline{m}) = \pi^L$; moreover, $\pi^H = \bigvee_{m' \in N^f_{l-1}} \sigma^{l-1}(m') \trianglerighteq \sigma^{l-1}(m)$; hence, $\langle \sigma^l(m) \rangle = \langle \sigma^{l-1}(m) \vee \pi^* \rangle = \langle \sigma^{l-1}(m) \rangle$. For $m \in N^f_{l-1}$ s.t. $m \ngeq f(l)$, $\langle \sigma^l(m) \rangle = \langle \sigma^{l-1}(m) \rangle$. Since by the inductive hypothesis, $\langle \sigma^{l-1}(m) \rangle = \beta(m)$, we have established that $\langle \sigma^l(m) \rangle = \beta(m)$ for all $m \in N^f_{l-1}$. We now need to show that σ^l is monotone. Consider any $m, m' \in N^f_{l-1}$ s.t. $m \geq m'$. There are three cases. First, suppose $m \geq m' \geq f(l)$. In that case, we know that $\sigma^l(m) = \sigma^{l-1}(m) \vee \pi^*$ and $\sigma^l(m') = \sigma^{l-1}(m') \vee \pi^*$. Since (by the inductive hypothesis) $\sigma^{l-1}(m) \trianglerighteq \sigma^{l-1}(m')$, we know that $\sigma^{l-1}(m) \vee \pi \trianglerighteq \sigma^{l-1}(m') \vee \pi$ for all π, and hence $\sigma^l(m) \trianglerighteq \sigma^l(m')$. The second case is where $m \geq f(l)$ and $m' \ngeq f(l)$. Then, $\sigma^l(m) = \sigma^{l-1}(m) \vee \pi^*$ and $\sigma^l(m') = \sigma^{l-1}(m')$. Since (by the inductive hypothesis) $\sigma^{l-1}(m) \trianglerighteq \sigma^{l-1}(m')$, we have that $\sigma^l(m) = \sigma^{l-1}(m) \vee \pi^* \trianglerighteq \sigma^{l-1}(m) \trianglerighteq \sigma^{l-1}(m') = \sigma^l(m')$. Finally, suppose that $m \ngeq f(l)$ and $m' \ngeq f(l)$. Then, $\sigma^l(m) = \sigma^{l-1}(m)$ and $\sigma^l(m') = \sigma^{l-1}(m')$. Since (by the inductive hypothesis) $\sigma^{l-1}(m) \trianglerighteq \sigma^{l-1}(m')$, we have that $\sigma^l(m) \trianglerighteq \sigma^l(m')$.

Finally, suppose $f(l) \in \mathit{Disconnected}\left(N^f_{l-1}\right)$. We assign an arbitrary signal $\sigma^l(f(l))$ to $f(l)$ such that $\langle \sigma^l(f(l)) \rangle = \beta(f(l))$, and we keep the signal allocation to nodes in N^f_{l-1} unchanged, i.e., $\sigma^l(m) = \sigma^{l-1}(m)$ for all $m \in N^f_{l-1}$. It is clear that σ^l induces appropriate beliefs (by the inductive hypothesis for $m \in N^f_{l-1}$ and by construction for $f(l)$). Since $f(l)$ is not comparable to any node in N^f_{l-1}, the fact that the signal allocation on N^f_{l-1} is monotone implies that the signal allocation on N^f_l is also monotone. This completes the proof. □

A.3 Proof of Proposition 2

Given a path P, directed or undirected, let N_P denote the set of nodes in P.

Lemma 4. *If H is K-universally constructible and H' is a closed subhierarchy of H, then H' is K-universally constructible.*

Proof of Lemma 4. Consider some (Ω, μ_0). Suppose H is a K-universally constructive hierarchy and $H' = (N', \geq)$ is a closed subhierarchy of H. Let β' be a monotone belief allocation on H'. We need to construct a monotone signal allocation on H' that induces β'. We define the belief allocation β on H as follows: (i) if $n \in N'$, let $\beta(n) = \beta'(n)$; (ii) if $n \notin N'$ and $\exists n' \in N'$ such that $n > n'$, let $\beta(n) = \overline{\tau}$; and (iii) if $n \notin N'$ and $\nexists n' \in N'$ such that $n > n'$, let $\beta(n) = \underline{\tau}$.

Claim 2. β is monotone on H.

Proof of Claim 2. Consider $n, n' \in N$ with $n \geq n'$. We need to show that $\beta(n) \succsim \beta(n')$. If n and n' are both in N', this follows from the fact that β' is monotone on H'. If neither n nor n' are in N', consider two cases. If $\exists n'' \in N'$ such that $n > n''$, then $\beta(n) = \overline{\tau} \succsim \beta(n')$. Otherwise, since $n \geq n'$ and $\nexists n'' \in N'$ such that $n > n''$, it must be that $\nexists n'' \in N'$ such that $n' > n''$, so $\beta(n) \succsim \beta(n') = \underline{\tau}$. If $n \notin N'$ and $n' \in N'$, $\beta(n) = \overline{\tau} \succsim \beta(n')$. Finally, if $n \in N'$ and $n' \notin N'$, it will suffice to show that there does not exist $n'' \in N'$ with $n' > n''$. Suppose to the contrary that there is such n''. Then, $n' \in Btw(n, n'')$. Since H' is closed and $n, n'' \in N'$, we have that $n' \in N'$ and we have reached a contradiction. ♢

Since β is monotone on H, and H is K-universally constructible, there exists a monotone signal allocation σ on H that induces β. Clearly, the restriction of σ to N' induces β' and is monotone on H'. □

The *between set of* (n', n'') is defined as

$$Btw\left(n', n''\right) = \left\{n \in N | n' \geq n \geq n''\right\}.$$

Clearly, the subhierarchy induced by the between set of any pair of nodes is closed. Moreover, if H' is closed, then N' contains $Btw(n', n'')$ for all $n', n'' \in N'$. We say a between set $Btw(n', n'')$ is *simple* if every $n \in Btw(n', n'') \setminus \{n', n''\}$ belongs to exactly one directed path in $G(H)$ from n' to n''. H' is a *minimal cyclic closed subhierarchy* (*MCC*) if it is cyclic, closed, and there is no cyclic and closed subhierarchy $H'' = (N'', \geq)$ with $N'' \subsetneq N'$. We say that a cycle in $\tilde{G}(H')$ is a *spanning cycle* if every node in H' is in the cycle.

Lemma 5. *If a subhierarchy $H' = (N', \geq)$ is an MCC, then either (i) N' is a simple between set, or (ii) every cycle in $\tilde{G}(H')$ is a spanning cycle.*

Proof of Lemma 5. Suppose H' is an MCC. We consider two cases. First, suppose there are two nodes $n', n'' \in N'$ such that $n' \geq n''$ and there are two distinct paths from n' to n'' in the directed graph $G(H)$. Then, the undirected graph of the subhierarchy $(Btw(n', n''), \geq)$ contains a cycle. We also know $(Btw(n', n''), \geq)$ is closed. Moreover, since H' is closed, we have that $Btw(n', n'') \subseteq N'$.

Hence, since H' is an MCC, we must have that $N' = Btw\,(n', n'')$. It remains to show that the between set N' is simple. Suppose to the contrary there is some node $\hat{n} \in N' \setminus \{n', n''\}$ such that $\hat{n}$ belongs to two distinct paths from n' to n'' in $G(H)$. Then, there are either two distinct directed paths from n' to $\hat{n}$ or two distinct directed paths from $\hat{n}$ to n''; thus, either $Btw\,(n', \hat{n})$ or $Btw\,(\hat{n}, n'')$ must be cyclic. Since both $Btw\,(n', \hat{n})$ and $Btw\,(\hat{n}, n'')$ are closed and strict subsets of N', H' must not be an MCC, so we have reached a contradiction. Thus, we have established that N' must be a simple between set.

Now consider the second case where for every $n', n'' \in N'$, there is at most one path from one node to the other in the directed graph $G(H)$. Since H' is an MCC, $\tilde{G}\,(H') = \left(N', \tilde{E}\right)$ contains a cycle $C = (n_0, \tilde{e}_0, ...n_{L-1}, \tilde{e}_{L-1}, n_L)$ where $L > 1$, $n_0 = n_L$. We will argue that N_C is closed. The fact that $N_C = N'$ will then follow directly from the hypothesis that H' is an MCC.

Let us then suppose that N_C is not closed, in order to reach a contradiction. Given a directed path $(n_0, e_0, ..., n_{L-1}, e_{L-1}, n_L)$, its *undirected analog* is the undirected path $(n_0, \tilde{e}_0, ...n_{L-1}, \tilde{e}_{L-1}, n_L)$ where $\tilde{e}_i = \{n_i, n_{i+1}\}$. Say that a directed path P *only contains edges in* C if every edge in the undirected analog of P is in C. A directed path P in $G(H)$ is an *external directed connection* (EDC) from n to n' if (i) P is a directed path from n to n'; (ii) $n, n' \in N_C$; and (iii) P does not only contain edges in C. Say that $(n, n') \in N_C$ are an *externally connected pair* (ECP) if there is an external directed connection from n to n' or from n' to n. An ECP (n_i, n_j) is said to be *minimally close* if for every $i \leq \underline{l} < \bar{l} \leq j$, $\left(n_{\underline{l}}, n_{\bar{l}}\right)$ is an ECP only if $\underline{l} = i$ and $\bar{l} = j$.

Claim 3. Given any two nodes $n, m \in N'$, if P is the unique directed path from n to m, then $N_P = Btw\,(n, m)$.

Proof of Claim 3. If there are two non-comparable nodes in $Btw\,(n, m)$, there would be two distinct directed paths from n to m. Hence, all nodes in $Btw\,(n, m)$ are comparable. Therefore, there is a directed path from n to m whose nodes are $Btw\,(n, m)$. Since there is a unique directed path from n to m, the set of nodes in P is $Btw\,(n, m)$. ♢

Claim 4. There exist i, j such that (n_i, n_j) is a minimally close ECP.

Proof of Claim 4. We know there is a pair of nodes in N_C that are an ECP. Otherwise, $(N_C, \geq)$ would be closed. Moreover, since L is finite, there is a pair of nodes in N_C that are a minimally close ECP, say (n_i, n_j). $\diamondsuit$

Let (n_i, n_j) be a minimally close ECP s.t. $\{n_i, n_{i+1}, ..., n_j\} \subsetneq N_C$. Let $\overline{n} = \max\{n_i, n_j\}$ and $\underline{n} = \min\{n_i, n_j\}$. Let P^e denote the external directed connection from $\overline{n}$ to $\underline{n}$. Let $\tilde{P}$ be the undirected path $(n_i, \tilde{e}_i, ..., \tilde{e}_{j-1}, n_j)$ from n_i to n_j. Let $\tilde{Q}$ denote the undirected path from n_i to n_j that "goes in the other direction" from $\tilde{P}$ in C, i.e., $\tilde{Q} = (n_i, \tilde{e}_{i-1}, n_{i-1}, ..., \tilde{e}_0, n_0, \tilde{e}_{L-1}, n_{L-1}, ..., \tilde{e}_j, n_j)$. Let $S = N_{\tilde{P}} \cup Btw(\overline{n}, \underline{n})$.

Claim 5. $(S, \geq)$ is cyclic.

Proof of Claim 5. It suffices to show there are two distinct undirected paths from n_i to n_j in $\tilde{G}((S, \geq))$. One path is $\tilde{P}$. The other path is the undirected analog of the external directed connection P^e. Since P^e is external, these two undirected paths must be distinct. $\diamondsuit$

Claim 6. S is closed.

Proof of Claim 6. Let $Y = \cup_{n,n' \in N_{\tilde{P}}} Btw(n, n')$. We will show that Y is closed and that $Y = S$.

First we show that Y is closed. Consider any $n', n'' \in Y$ and $n \in Btw(n', n'')$. By definition of Y, $n' \in Btw(n_1, n_2)$ and $n'' \in Btw(n_3, n_4)$, where $n_l \in N_{\tilde{P}}$ for $l = 1, 2, 3, 4$. Hence, $n_1 \geq n' \geq n \geq n'' \geq n_4$ and thus $n \in Btw(n_1, n_4) \subseteq Y$.

It remains to show that $S = Y$. Given $n \in N_{\tilde{P}}$, $n \in Btw(n, n) \subseteq Y$. Moreover, $Btw(\overline{n}, \underline{n}) \subseteq Y$. Hence, $S \subseteq Y$.

Now, consider some $n \in Y$. We need to show that $n \in S$. If $n \in N_{\tilde{P}}$, then we are done. Otherwise, $n \notin N_{\tilde{P}}$. We know $n \in Btw(n', n'')$ for some $n', n'' \in N_{\tilde{P}}$. If $(n', n'') = (\overline{n}, \underline{n})$, $n \in Btw(\overline{n}, \underline{n}) \subseteq S$. Suppose instead that $(n', n'') \neq (\overline{n}, \underline{n})$. We will reach a contradiction. Let P denote the directed path from n' to n'' whose nodes include n. Because $(\overline{n}, \underline{n})$ is a minimally close ECP, path P must only include edges in C. Since $n \notin N_{\tilde{P}}$, the nodes in path P cannot be a subset of $N_{\tilde{P}}$. Thus, the nodes in P contain the nodes in $\tilde{Q}$, including $\overline{n}$ and $\underline{n}$. The sequence of nodes and edges in P between $\overline{n}$ and $\underline{n}$ is a directed path between those nodes, and thus is equal to P^e (by uniqueness of the directed path). Since P^e contains an edge which is not in C, we have contradicted the hypothesis that P only contains edges in C. Thus, we have established that $Y \subseteq S$. $\diamondsuit$

We have established that $(S, \geq)$ is cyclic and closed and that $S \subseteq N'$. Since H' is an MCC, it must be that $S = N'$. But since $S = N_{\tilde{P}} \cup Btw\left(\overline{n}, \underline{n}\right)$, it must be that $N_{\tilde{Q}} \subseteq Btw\left(\overline{n}, \underline{n}\right)$. All nodes in a between set are comparable, by the assumption of unique directed paths, and so all nodes in $N_{\tilde{Q}}$ are comparable. Hence, $\tilde{Q}$ must be the undirected analogue of P^e. This contradicts the hypothesis that $\left(\overline{n}, \underline{n}\right)$ is an ECP. $\square$

Lemma 6. *Given a cyclic subhierarchy $H' = (N', \geq)$, if every cycle in $\tilde{G}(H')$ is a spanning cycle, then for any pair of nodes $n, m \in N'$, there exist two undirected paths from n to m such that the union of the nodes in the two paths is N' and the intersection of the nodes in the two paths is $\{n, m\}$.*

Proof of Lemma 6. Since there exists a spanning cycle $\tilde{G}(H')$, for any pair of nodes $n, m \in N'$, there exist two undirected paths $P = \left(n, \tilde{e}_0^P, n_1^P ..., n_{L^P-1}^P, \tilde{e}_{L^P-1}^P, m\right)$ and $Q = \left(n, \tilde{e}_0^Q, n_1^Q ..., n_{L^Q-1}^Q, \tilde{e}_{L^Q-1}^Q, m\right)$ such that $N' = N_P \cup N_Q$ and $\tilde{E}_P \cap \tilde{E}_Q = \emptyset$, where $\tilde{E}_P = \left\{\tilde{e}_0^P, ..., \tilde{e}_{L^P-1}^P\right\}$ and $\tilde{E}_Q = \left\{\tilde{e}_0^Q, ..., \tilde{e}_{L^Q-1}^Q\right\}$. We need to show that $N_P \cap N_Q = \{n, m\}$. Suppose to the contrary that there exists $\hat{n} \in N_P \cap N_Q$ with $\hat{n} \notin \{n, m\}$. We know there exist $l^P \in \{1, ..., L^P - 1\}$ and $l^Q \in \{1, ..., L^Q - 1\}$ such that $\hat{n} = n_{l^P}^P = n_{l^Q}^Q$. Now, consider the undirected path $\left(n, \tilde{e}_0^P, n_1^P, ..., \tilde{e}_{l^P-1}^P, n_{l^P}^P, \tilde{e}_{l^Q-1}^Q, n_{l^Q-1}^Q, \tilde{e}_{l^Q-2}^Q, ..., n\right)$. Since $\tilde{E}_p \cap \tilde{E}_q = \emptyset$, this is a well-defined path. But it is a cycle that is not spanning and thus we have reached a contradiction. $\square$

Next, let us say that a signal π is *simple* if for every $s \in \pi$, $p(s) > 0$. A signal allocation is simple if every signal in its range is simple.

Lemma 7. *For any hierarchy H, a monotone belief allocation β is induced by some monotone signal allocation σ if and only if there exists a simple monotone signal allocation σ' that also induces β.*

Proof of Lemma 7. Let

$$\pi = \vee_{n \in N} \sigma(n)$$

be the join of all of the signals (which is clearly finite), and let

$$\hat{s} = \cup \left\{s' \in \pi | p\left(s'\right) = 0\right\}.$$

We define a new signal allocation as follows. Pick some $\tilde{s} \in \pi$ such that $p(\tilde{s}) > 0$, and define the mapping $f : 2^{\Omega} \to 2^{\Omega}$:

$$f(s) = \begin{cases} s \cup \hat{s} & \text{if } \tilde{s} \subseteq s; \\ s \setminus \hat{s} & \text{otherwise,} \end{cases}$$

and define the new signal allocation

$$\sigma'(n') = \left\{ f(s) | s \in \sigma(n') \right\}.$$

Then clearly for each n, all we have done is remove a zero measure set from positive measure sets and added it to the (unique) positive measure set that contains $\tilde{s}$, so that $\sigma'(n')$ is simple. Moreover, it is clear that by moving around a mass of zero, we have not changed the induced distribution of beliefs at any node. The new signal allocation is monotone as well: if $n < n'$, then every $s' \in \sigma(n')$ is contained in some $s \in \sigma(n)$. If $\tilde{s} \subseteq s'$, then clearly $\tilde{s} \subseteq s$, so $f(s') = s' \cup \hat{s} \subseteq s \cup \hat{s} = f(s)$. If $\tilde{s} \nsubseteq s'$, then $f(s') = s' \setminus \hat{s} \subseteq s \setminus \hat{s} \subseteq f(s)$ (the inclusion may be strict if $\tilde{s} \subseteq s \setminus s'$). $\square$

Suppose a subhierarchy H' has nodes

$$N' = \{\overline{n}\} \cup \{\underline{n}\} \cup_{l=1}^{L} \left\{ n_k^l \right\}_{k=1}^{K_l},$$

with $L \geq 2$, such that for all $l = 1, \ldots, L$, (i) $K_l \geq 1$, (ii) $\overline{n} > n_1^l$, (iii) $n_{K_l}^l > \underline{n}$, (iv) n_k^l and n_{k+1}^l are comparable for every $k = 1, \ldots, K_l - 1$, and (v) n_k^l is not comparable to $n_{k'}^{l'}$ if $l \neq l'$. In this case, we say that H' is a *union of non-comparable paths (UNP).*

Lemma 8. *If H' is a UNP, then H' is not 2-universally constructible.*

Proof of Lemma 8. First, we establish the following martingale property:

Claim 7. If $\pi' \trianglerighteq \pi$, then for all $s \in \pi$ with $p(s) > 0$,

$$\mu_s = \sum_{s' \in \pi'} \mu_{s'} \frac{p(s \cap s')}{p(s)}. \tag{1}$$

Proof of Claim 7. This follows from the following chain of equalities: for all $\omega \in \Omega$,

$$\begin{aligned}\mu_s &= \frac{p(s|\omega)\,\mu_0(\omega)}{p(s)} \\ &= \sum_{\{s' \in \pi' | s' \subseteq s\}} \frac{p(s'|\omega)\,\mu_0(\omega)}{p(s)} \\ &= \sum_{\{s' \in \pi' | s' \subseteq s\}} \mu_{s'}(\omega)\,\frac{p(s')}{p(s)} \\ &= \sum_{s' \in \pi'} \mu_{s'}(\omega)\,\mathbb{I}_{s' \subseteq s} \frac{p(s')}{p(s)},\end{aligned}$$

which is equal to the right-hand side of (1), since $\mathbb{I}_{s' \subseteq s} p(s') = p(s' \cap s)$ due to the fact that $\pi' \trianglerighteq \pi$. ♢

To prove the lemma, we will construct a monotone belief allocation β, based on the belief allocation on the diamond in Section 3, and show that β cannot be induced by any monotone signal allocation. Let $\beta(\overline{n}) = \tau_{\otimes}$; $\beta(\underline{n}) = \tau_{\bigcirc}$; for $k = 1, \ldots, K_1$, let $\beta(n_k^1) = \tau_{\obslash}$; and for all $l = 2, \ldots, L$ and $k = 1, \ldots, K_l$, let $\beta(n_k^l) = \tau_{\oslash}$. Recall that $\tau_{\otimes}$ has support $\{0, \frac{1}{2}, 1\}$, $\tau_{\bigcirc}$ has support $\{\frac{1}{6}, \frac{5}{6}\}$, $\tau_{\obslash}$ has support $\{0, \frac{2}{3}, \frac{5}{6}\}$, and $\tau_{\oslash}$ has support $\{\frac{1}{6}, \frac{1}{3}, 1\}$. By the argument in footnote 8, β is monotone.

Now, by Lemma 7, H' is 2-universally constructible only if there exists a simple signal allocation σ that is monotone with respect to H' and induces β. Fix such a signal allocation. We will need a preliminary result. Given a signal π and $\mu \in \text{supp}\langle\pi\rangle$, define the set $s(\pi, \mu) = \cup\{s' \in \pi | p(s') > 0 \text{ and } \mu_{s'} = \mu\}$. This is the set of all outcomes in $\Omega \times [0,1]$ under which the belief μ is generated by the signal π. Note that if π is a simple signal, then the sets $\{s(\pi, \mu) | \mu \in \text{supp}\langle\pi\rangle\}$ are themselves a signal, i.e., a partition of $\Omega \times [0,1]$, that is a coarsening of π. Note also that if π and π' are simple and $\pi \trianglerighteq \pi'$, then for every $s' \in \pi'$, $s' = \cup\{s \in \pi | p(s \cap s') > 0\}$.

Claim 8. For all $l = 1, \ldots, L$ and for all $\mu \in \text{supp}\langle\beta(n_1^l)\rangle$, $s(\sigma(n_k^l), \mu) = s(\sigma(n_1^l), \mu)$ for all $k = 1, \ldots, K_l$.

Proof of Claim 8. The proof is by induction on k, with the base case of $k = 1$ being a hypothesis of the claim. For the inductive step, suppose that $s(\sigma(n_k^l), \mu) = s(\sigma(n_1^l), \mu)$. By assumption,

n_k^l and n_{k+1}^l are comparable. Assume for now that $n_k^l > n_{k+1}^l$. By (1), for every $s \in \sigma\left(n_{k+1}^l\right)$,

$$\mu_s = \sum_{s' \in \sigma\left(n_k^l\right)} \mu_{s'} \frac{p\left(s' \cap s\right)}{p\left(s\right)}.$$

Now suppose that $\mu_{s'} \neq \mu_s$ for some s' with $p\left(s' \cap s\right) > 0$. Then we would conclude that $\langle \sigma\left(n_k^l\right) \rangle$ is a strict mean-preserving spread of $\langle \sigma\left(n_{k+1}^l\right) \rangle$, which contradicts the hypothesis that both signals induce the same beliefs. Thus, it must be that for all $s' \in \sigma\left(n_k^l\right)$ with $p\left(s' \cap s\right) > 0$, $\mu_{s'} = \mu_s$. As a result,

$$\begin{aligned} s\left(\sigma\left(n_{k+1}^l\right), \mu\right) &= \cup_{\left\{s \in \sigma\left(n_{k+1}^l\right) | \mu_s = \mu\right\}} s \\ &= \cup_{\left\{s' \in \sigma\left(n_k^l\right) | \exists s \in \sigma\left(n_{k+1}^l\right) \text{ s.t. } p(s' \cap s) > 0 \text{ and } \mu_s = \mu\right\}} s \\ &\subseteq \cup_{\left\{s' \in \sigma\left(n_k^l\right) | \mu_{s'} = \mu\right\}} s \\ &= s\left(\sigma\left(n_k^l\right), \mu\right). \end{aligned}$$

But since both sets have the same measure and the signals are simple, we conclude that $s\left(\sigma\left(n_{k+1}^l\right), \mu\right) = s\left(\sigma\left(n_k^l\right), \mu\right) = s\left(\sigma\left(n_1^l\right), \mu\right)$ as desired. The case where $n_{k+1}^l > n_k^l$ is analogous and is omitted. $\diamondsuit$

We will now reach a contradiction by first arguing that $p\left(s\left(\sigma\left(\overline{n}\right), 0\right) \cap s\left(\sigma\left(\underline{n}\right), 5/6\right)\right) = 0$ and then arguing that $p\left(s\left(\sigma\left(\overline{n}\right), 0\right) \cap s\left(\sigma\left(\underline{n}\right), 5/6\right)\right) = 1/24$, i.e., that $Pr(\tilde{\mu}_{\sigma(\overline{n})} = 0 \;\&\; \tilde{\mu}_{\sigma(\underline{n})}=5/6)$ is equal to both 0 and 1/24.

Step 1: $p\left(s\left(\sigma\left(\overline{n}\right), 0\right) \cap s\left(\sigma\left(\underline{n}\right), 5/6\right)\right) = 0$. First, we argue that $s\left(\sigma\left(\overline{n}\right), 0\right) = s\left(\sigma\left(n_1^1\right), 0\right)$. The reason is that for any $s \in \sigma\left(n_1^1\right)$ with $\mu_s = 0$ and $s' \in \sigma\left(\overline{n}\right)$, the martingale property implies that $p\left(s \cap s'\right) > 0$ only if $\mu_{s'} = 0$. As a result, $s\left(\sigma\left(n_1^1\right), 0\right) \subseteq s\left(\sigma\left(\overline{n}\right), 0\right)$. But since these two sets have the same measure, which is 3/8, and since $\sigma\left(\overline{n}\right)$ is simple, the two sets must be equal. Next, a similar argument indicates that $s\left(\sigma\left(n_{K_1}^1\right), 5/6\right) = s\left(\sigma\left(\underline{n}\right), 5/6\right)$. This implies that $s\left(\sigma\left(n_{K_1}^1\right), 0\right) \subseteq s\left(\sigma\left(\underline{n}\right), 1/6\right)$. Finally, Claim 8 implies that $s\left(\sigma\left(n_1^1\right), 0\right) = s\left(\sigma\left(n_{K_1}^1\right), 0\right)$.

Putting it all together,

$$
\begin{aligned}
s\left(\sigma\left(\overline{n}\right),0\right) &= s\left(\sigma\left(n_1^1\right),0\right)\\
&= s\left(\sigma\left(n_{K_1}^1\right),0\right)\\
&\subseteq s\left(\sigma\left(\underline{n}\right),1/6\right),
\end{aligned}
$$

so that $s\left(\sigma\left(\overline{n}\right),0\right)\cap s\left(\sigma\left(\underline{n}\right),5/6\right)=\emptyset$.

Step 2: $p\left(s\left(\sigma\left(\overline{n}\right),0\right)\cap s\left(\sigma\left(\underline{n}\right),5/6\right)\right)=1/24$. Similar arguments as above indicate that $s\left(\sigma\left(n_{K_2}^2\right),1/6\right)=s\left(\sigma\left(\underline{n}\right),1/6\right)$ and that $s\left(\sigma\left(n_1^2\right),1\right)=s\left(\sigma\left(\overline{n}\right),1\right)$. The fact that $s\left(\sigma\left(n_{K_2}^2\right),1/6\right)=s\left(\sigma\left(\underline{n}\right),1/6\right)$ means that $s\left(\sigma\left(n_{K_2}^2\right),1/3\right)\cup s\left(\sigma\left(n_{K_2}^2\right),1\right)=s\left(\sigma\left(\underline{n}\right),5/6\right)$. Next, the fact that $s\left(\sigma\left(n_1^2\right),1\right)=s\left(\sigma\left(\overline{n}\right),1\right)$ means that $s\left(\sigma\left(n_1^2\right),1/3\right)\subseteq s\left(\sigma\left(\overline{n}\right),0\right)\cup s\left(\sigma\left(\overline{n}\right),1/2\right)$. The martingale property then implies $p\left(s\left(\sigma\left(\overline{n}\right),0\right)\cap s\left(\sigma\left(n_1^2\right),1/3\right)\right)=1/24$ (with $p\left(s\left(\sigma\left(\overline{n}\right),1/2\right)\cap s\left(\sigma\left(n_1^2\right),1/3\right)\right)=1/12$ and $p\left(s\left(\sigma\left(n_1^2\right),1/3\right)\right)=1/8$). Finally, Claim 8 implies that $s\left(\sigma\left(n_1^2\right),1/3\right)=s\left(\sigma\left(n_{K_2}^2\right),1/3\right)$ and that $s\left(\sigma\left(n_1^2\right),1\right)=s\left(\sigma\left(n_{K_2}^2\right),1\right)$. Putting it all together,

$$
\begin{aligned}
s\left(\sigma\left(\overline{n}\right),0\right)\cap s\left(\sigma\left(\underline{n}\right),5/6\right) &= s\left(\sigma\left(\overline{n}\right),0\right)\cap\left(s\left(\sigma\left(n_{K_2}^2\right),1/3\right)\cup s\left(\sigma\left(n_{K_2}^2\right),1\right)\right)\\
&= s\left(\sigma\left(\overline{n}\right),0\right)\cap\left(s\left(\sigma\left(n_1^2\right),1/3\right)\cup s\left(\sigma\left(n_1^2\right),1\right)\right)\\
&= s\left(\sigma\left(\overline{n}\right),0\right)\cap\left(s\left(\sigma\left(n_1^2\right),1/3\right)\cup s\left(\sigma\left(\overline{n}\right),1\right)\right)\\
&= s\left(\sigma\left(\overline{n}\right),0\right)\cap s\left(\sigma\left(n_1^2\right),1/3\right),
\end{aligned}
$$

yielding $p\left(s\left(\sigma\left(\overline{n}\right),0\right)\cap s\left(\sigma\left(\underline{n}\right),5/6\right)\right)=p\left(s\left(\sigma\left(\overline{n}\right),0\right)\cap s\left(\sigma\left(n_1^2\right),1/3\right)\right)=1/24$, as desired. □

Lemma 9. *Given a cyclic subhierarchy $H'=(N',\geq)$, if N' is a simple between set, then H' is not 2-universally constructible.*

Proof of Lemma 9. A simple between set is a UNP, where $\overline{n}$ and $\underline{n}$ are the maximal and minimal elements of N', respectively; L is the number of distinct directed paths between $\overline{n}$ and $\underline{n}$; and the sequences $\{n_k^l\}_{k=1}^{K_l}$ are the nodes in these paths (excluding the end nodes of course). From cyclicity, there must be at least two such paths so $L\geq 2$, with $K_l\geq 1$ for each l. Lemma 8 then implies that H' is not 2-universally constructible. □

Lemma 10. *Given a cyclic subhierarchy $H' = (N', \geq)$, if H' is not a crown subhierarchy, and every cycle in $\tilde{G}(H')$ is a spanning cycle, then H' is not 2-universally constructible.*

Proof of Lemma 10. Suppose that H' is not a crown subhierarchy and that every cycle in $\tilde{G}(H')$ is a spanning cycle. We say an element $\overline{n}$ of N' is maximal if there is no $n \in N'$ such that $n > \overline{n}$. We say an element $\underline{n}$ of N' is minimal if there is no $n \in N'$ such that $\underline{n} > n$.

Claim 9. There exist $\overline{n}$ and $\underline{n}$ in N' that are maximal and minimal, respectively, such that $\overline{n}$ does not cover $\underline{n}$.

Proof of Claim 9. Suppose toward contradiction that every maximal element covers every minimal element. Hence, $\tilde{G}(H')$ is a complete bipartite graph of maximal and minimal elements. If there were only one maximal element or only one minimal element, there could not be a cycle. So, there must be at least two of each. Take any $N'' \subseteq N'$ consisting of exactly two maximal and two minimal elements, and let $H'' = (N'', \geq)$. H'' is a crown subhierarchy and therefore is cyclic. Since every cycle in $\tilde{G}(H')$ is a spanning cycle, we must have $N' = N''$, and thus H' is a crown subhierarchy. Hence, we have reached a contradiction. ♢

By Lemma 6, there are two undirected paths P and Q in $\tilde{G}(H')$ from $\overline{n}$ to $\underline{n}$ such that the union of the nodes in the two paths is N' and the intersection of the nodes in the two paths is $\{\overline{n}, \underline{n}\}$. As a result, H' satisfies the hypotheses of Lemma 8, where $\overline{n}$ and $\underline{n}$ are as stated, $L = 2$, $\{n_k^1\}_{k=1}^{K_1}$ are the interior nodes in P, and $\{n_k^2\}_{k=1}^{K_2}$ are the interior nodes in Q. Lemma 8 therefore implies that H' is not 2-universally constructible. □

Proof of Proposition 2. Suppose $\tilde{G}(H)$ is not a forest, i.e., it contains a cycle. Therefore, H is closed and cyclic, and—since N is finite—H contains a subhierarchy $H' = (N', \geq)$ that is an MCC. By Lemma 5, either (i) N' is a simple between set, or (ii) every cycle in $\tilde{G}(H')$ is a spanning cycle. Consider case (i). Then, Lemma 9 implies that H' is not 2-universally constructible. Now consider case (ii). If H' is not a crown subhierarchy, then Lemma 10 implies that is not 2-universally constructible. If H' is a crown subhierarchy, Lemma 3 implies it is not 3-universally constructible. Hence, if $K \geq 3$, H' is not K-universally constructible. Therefore, by Lemma 4, H is not K-universally constructible for $K \geq 3$. □