

INFORMATION HIERARCHIES

BENJAMIN BROOKS

Department of Economics, University of Chicago

ALEXANDER FRANKEL

Booth School of Business, University of Chicago

EMIR KAMENICA

Booth School of Business, University of Chicago

If experiment A is Blackwell more informative than experiment B , it is always possible that A and B are induced by signals A' and B' such that A' is a refinement of B' , that is, A' entails observing B' plus some additional information. We first show that this result does not extend beyond pairs of experiments: There exist collections of experiments that cannot be induced by a collection of signals so that whenever two experiments are Blackwell ordered, the associated signals are refinement ordered. In other words, sometimes it is impossible for more informed agents to know everything that less informed agents know. More broadly, define an *information hierarchy* to be a partially ordered set that ranks experiments in terms of informativeness. Is it the case that for any choice of experiments indexed on the hierarchy such that higher experiments are Blackwell more informative, there are signals that induce these experiments with higher signals being refinements of lower signals? We show that the answer is affirmative if and only if the undirected graph of the information hierarchy is a forest.

KEYWORDS: Information, beliefs, signals, Blackwell order, networks.

1. INTRODUCTION

THERE ARE TWO DISTINCT THINGS we might mean when we say that A is more informed than B . One is that A 's information about some state of the world is more accurate than B 's. The other is that A knows everything that B knows; that is, A observes what B observes, plus something else. The first notion is typically formalized by representing a source of information as an experiment, that is, a distribution of posterior beliefs (Blackwell (1953)).¹ We then say that A is Blackwell more informed than B if A 's beliefs are a mean-preserving spread of B 's beliefs. The second notion is formalized by representing a source of information as a partition of some expanded state space $\Omega \times X$ (Green and Stokey (1978)). This formulation distinguishes payoff-relevant states (Ω) from the realizations of signals conditional on those states (X). We then say that A knows everything that B knows if A 's partition is a refinement of B 's partition.

These two formalisms are closely related. Every signal, that is, partition of the expanded state space, induces an experiment. If one signal is a refinement of another, then the experiment it induces is Blackwell more informative. Moreover, Green and Stokey

Benjamin Brooks: babrooks@uchicago.edu

Alexander Frankel: afrankel@uchicago.edu

Emir Kamenica: emir@uchicago.edu

We thank Ozan Candogan, Eddie Dekel, Simone Galperti, Stephen Morris, Xiaosheng Mu, Ludvig Sinander, Alex Smolin, and Bruno Strulovici for helpful comments. We thank Oguz Bayraktar and Cheryl Wu for research assistance. Kamenica thanks the Neubauer Family Faculty Fellowship for research support.

¹This definition is equivalent to the prior-free formulation of an experiment as a map from the state to distributions over signal realizations.

(1978) established a partial converse: given two Blackwell-ordered experiments, there exist refinement-ordered signals that induce those experiments. In other words, if A 's information is more accurate than B 's, it is always possible that A knows everything that B knows.

A natural question is whether Green and Stokey's result extends beyond pairs of experiments. Suppose we have some collection of experiments. Is it always the case that we can construct a collection of signals that induce those experiments, and whenever two experiments are Blackwell ordered, the two corresponding signals are refinement ordered?

We show that the answer is negative. We construct an example where A is Blackwell more informed than B and C , who are in turn more informed than D , and yet it cannot be that A knows everything that B and C know and that all three know everything that D knows.

To examine this issue more generally, we introduce the notion of an *information hierarchy*, which is simply a partially ordered set. We consider experiment allocations that assign an experiment to each element of the hierarchy, with the property that a higher element's experiment is Blackwell more informative than a lower element's experiment. We say an experiment allocation is constructible if it is possible to assign a signal to every element of the hierarchy so that each signal induces the corresponding experiment, and a higher element's signal is a refinement of a lower element's signal. If every monotone experiment allocation is constructible, we say that the information hierarchy is universally constructible.

Our main theorem characterizes the set of information hierarchies that are universally constructible. A partially ordered set (and thus an 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 establish that an information hierarchy is universally constructible if and only if this graph is a forest.

We also establish a generalization of this result. We consider a class of other economically meaningful orders on signals, and we show that, if we replace refinement with any one of these orders in our notion of constructibility, our theorem continues to hold.

We discuss two applications of our theorem. First, suppose there is an agent who observes information from several sources and an econometrician who does not know the data generating process behind these sources. We analyze the conditions under which the econometrician can rationalize the agent's reactions to this information. In particular, we show that two necessary conditions, Bayes plausibility and Blackwell monotonicity,³ do not suffice for rationalizability when there are more than two sources of information. Second, we study information design in organizations under the constraint that managers must have access to the information of their subordinates.

The rest of this paper proceeds as follows. Section 2 establishes basic definitions. Section 3 presents motivating examples. Section 4 contains the remaining definitions and our main result. Section 5 discusses extensions. Section 6 provides two applications. We discuss related literature in probability theory and information economics in Section 7. Section 8 is a conclusion.

²The standard representation of a partially ordered set as a directed graph encodes the partial order by placing an edge from n to n' if n covers n' , that is, 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.

³Bayes plausibility requires that the agent's average belief does not vary across the sets of observed information sources. Blackwell monotonicity requires that when an agent observes additional sources, her beliefs become more dispersed.

2. STATES, SIGNALS, AND EXPERIMENTS

There is 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. Thus, the conditional probability of s given ω is $p(s|\omega) = \lambda(\{x | (\omega, x) \in s\})$ where $\lambda(\cdot)$ denotes the Lebesgue measure. Importantly, if we know that one agent observed some signal, say π_A , and another agent observed some signal π_B , we can deduce not only how informed each agent is about ω , but also the joint distribution of their signal realizations (and thus their higher-order beliefs).

In contrast, a (Blackwell) experiment only specifies how informed an agent is about ω , without specifying her information about the beliefs of other agents. The standard definition of an experiment is a map from Ω to distributions over signal realizations, but for ease of exposition we identify each experiment with the distribution of beliefs it induces. Thus, we define an *experiment*, denoted by τ , as an element of $\Delta(\Delta(\Omega))$ that has finite support and satisfies $\mathbb{E}_\tau[\mu] = \mu_0$.

We write $\pi \supseteq \pi'$ if π is a *refinement* of π' , that is, every element of π is a subset of some element of π' .⁴ If $\pi \supseteq \pi'$, an agent who observes π has access to all the information available to the agent who observes π' (and thus knows those agent's beliefs). We write $\tau \succcurlyeq \tau'$ if τ is *Blackwell more informative* than τ' , that is, τ is a mean-preserving spread of τ' . If $\tau \succcurlyeq \tau'$, an agent who observes τ obtains a higher payoff than the agent who observes τ' in any decision problem.

We denote the set of all signals by Π . The pair (Π, \supseteq) is a lattice and we let \vee denote the join, that is, $\pi \vee \pi'$ is the coarsest refinement of both π and π' . Note that $\pi \vee \pi'$ is the signal that is equivalent to observing both π and π' .

Each signal π induces an experiment, denoted $\langle \pi \rangle$, according to

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

where $p(s)$ is the unconditional probability of signal realization s and μ_s is the posterior belief about ω conditional on s .⁵ Moreover, given signal π , we let $\tilde{\mu}_\pi$ denote the belief-valued random variable on $\Omega \times [0, 1]$ that is the posterior induced by the observation of the signal realization from π . Note that the refinement order implies a belief-martingale property: if $\pi \supseteq \pi'$, then $\mathbb{E}[\tilde{\mu}_\pi | \tilde{\mu}_{\pi'}] = \tilde{\mu}_{\pi'}$.⁶ In other words, observing additional information cannot change one's belief on average.

3. ILLUSTRATIVE EXAMPLES

There are four experiments, τ_A, τ_B, τ_C , and τ_D , some of which might be Blackwell more informative than others. We wish to know whether there exist signals π_i for $i \in \{A, B, C, D\}$ such that

(C1) $\langle \pi_i \rangle = \tau_i$ and

⁴We then also say π' is *coarser* than π .

⁵Note that $p(s) = \sum_{\omega \in \Omega} p(s|\omega)\mu_0(\omega)$. For any s with $p(s) > 0$, we have $\mu_s(\omega) = \frac{p(s|\omega)\mu_0(\omega)}{p(s)}$. For those s with $p(s) = 0$, we can set μ_s to be any arbitrary belief.

⁶Throughout the paper, when we say two random variables are equal, we mean almost surely.

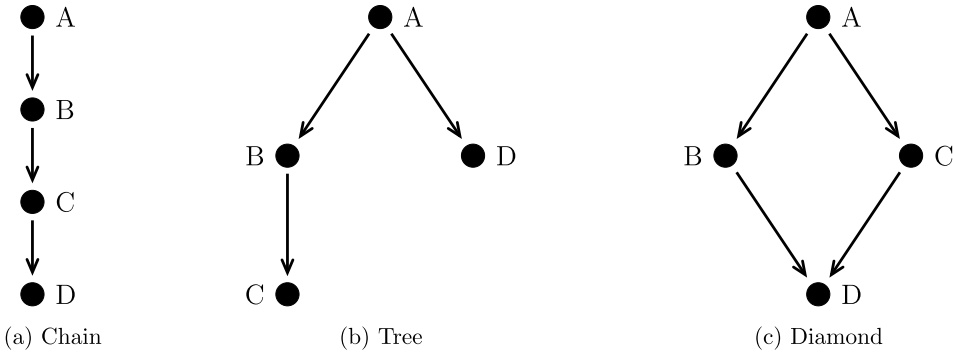


FIGURE 1.—Example hierarchies.

(C2) $\tau_i \succsim \tau_j$ implies $\pi_i \supseteq \pi_j$.

As we discussed in the [Introduction](#), the answer to this question will shed light on (i) whether Green and Stokey’s (1978) result on the equivalence of representing information as experiments versus signals generalizes beyond pairwise comparisons, and (ii) whether it is always possible that a more informed agent knows what all less informed agents know.

As we will see, for some Blackwell rankings of the experiments, the answer is yes, no matter what the the exact experiments are. For other rankings, this is not the case; there exist experiments for which there do not exist signals satisfying (C1) and (C2). We illustrate these results through three examples, depicted in Figure 1.

Chain

Suppose $\tau_A \succsim \tau_B \succsim \tau_C \succsim \tau_D$. When experiments are ordered in this way, it is always possible to construct signals that satisfy (C1) and (C2), regardless of the particular experiments. This follows from Green and Stokey (1978): For any π' and τ with $\tau \succsim \langle \pi' \rangle$, there exists a π such that (i) $\langle \pi \rangle = \tau$ and (ii) $\pi \supseteq \pi'$.⁷ In other words, take any signal that induces some experiment; there is a refinement of this signal that induces any particular more-informative experiment.

Now, consider any π_D that induces τ_D . Since $\tau_C \succsim \tau_D$, there exists some signal $\pi_C \supseteq \pi_D$ that induces τ_C . Similarly, there is a $\pi_B \supseteq \pi_C$ that induces τ_B . Finally, there is a $\pi_A \supseteq \pi_B$ that induces τ_A . Hence, we constructed signals that satisfy (C1) and (C2).

Tree

Next, suppose $\tau_A \succsim \tau_B \succsim \tau_C$, $\tau_A \succsim \tau_D$, and τ_D is not comparable with τ_B or τ_C . In this case, it is again always possible to construct signals so that $\langle \pi_i \rangle = \tau_i$ and $\tau_i \succsim \tau_j$ implies $\pi_i \supseteq \pi_j$. The argument, however, is more subtle than for the chain. To see the issue, suppose we first go “up the tree” and proceed as in the previous example: we construct π_D and then $\pi_A \supseteq \pi_D$ that induce τ_D and τ_A . Now, to assign a signal π_B , we have to go “down the tree.” But there is no guarantee that, given the π_A we constructed, there exists a π_B that is coarser than π_A and induces τ_B .

⁷This result first appears as Theorem 1 in Green and Stokey (1978). It is stated and proved using the notation in our paper by Gentzkow and Kamenica (2017).




FIGURE 2.—The diamond hierarchy is not universally constructible.

We establish a result (Lemma 1), however, that tells us we can construct π_B that induces τ_B and has the property that $\langle \pi_A \vee \pi_B \rangle = \langle \pi_A \rangle = \tau_A$. (To do so, we build a π_B whose realizations are independent of the state given the realization of π_A .) Then, we replace the “provisional” π_A with $\pi_A \vee \pi_B$ (which is of course also finer than π_D). We then apply a similar procedure to assign π_C and update the provisional π_A and π_B . Hence, we constructed signals that satisfy (C1) and (C2). This algorithm is discussed in greater detail in Section 4.

Diamond

Finally, suppose $\tau_A \succsim \tau_B \succsim \tau_D$, $\tau_A \succsim \tau_C \succsim \tau_D$, and τ_B and τ_C are not (necessarily) comparable. We now give an example of four experiments ordered in this manner such that there do not exist signals that satisfy (C1) and (C2).⁸

Let $\Omega = \{0, 1\}$ with a uniform prior. Since the state space is binary, we associate each belief with $\Pr(\omega = 1) \in [0, 1]$. The rows in Figure 2 depict the four experiments, with the number of dots in a circle proportional to the probability mass on that belief. It is easy to see that these four experiments satisfy the given Blackwell ordering, with arrows (solid or dashed) showing the unique way to obtain a higher-ranked distribution as a mean-preserving spread of a lower-ranked distribution.

Now, to obtain a contradiction, suppose that the signals π_A , π_B , π_C , and π_D satisfy (C1) and (C2). First, note that these signals do not merely induce the respective marginal distributions of beliefs (i.e., the experiments), but they also determine the joint distribution of these beliefs. In particular, as noted earlier, whenever $\pi \succeq \pi'$, we must have $\mathbb{E}[\tilde{\mu}_\pi | \tilde{\mu}_{\pi'}] = \tilde{\mu}_{\pi'}$.

We use this fact to show that when $\tilde{\mu}_{\pi_D} = \frac{1}{6}$, (i) $\pi_A \succeq \pi_B \succeq \pi_D$ implies that $\tilde{\mu}_{\pi_A} \neq 1$, while (ii) $\pi_A \succeq \pi_C \succeq \pi_D$ implies that $\tilde{\mu}_{\pi_A} = 1$ with a strictly positive probability. To see (i), note that when $\tilde{\mu}_{\pi_D} = \frac{1}{6}$, we must have $\tilde{\mu}_{\pi_B} = \frac{1}{6}$ and thus $\tilde{\mu}_A \neq 1$ (since $\tilde{\mu}_{\pi_A} = 1$ if and only if $\tilde{\mu}_{\pi_B} = 1$). To see (ii), note that when $\tilde{\mu}_{\pi_D} = \frac{1}{6}$, there is a strictly positive probability that $\tilde{\mu}_{\pi_C} = \frac{2}{3}$ (since $\tilde{\mu}_{\pi_C} = \frac{5}{6}$ if and only if $\tilde{\mu}_{\pi_D} = \frac{5}{6}$); and when $\tilde{\mu}_{\pi_C} = \frac{2}{3}$, there is a strictly positive probability that $\tilde{\mu}_{\pi_A} = 1$ (since all of the other beliefs in the support of $\tilde{\mu}_{\pi_A}$ are strictly lower than $\frac{2}{3}$). Thus, we have reached a contradiction.

⁸The mathematical details of this example are presented in the proof of Lemma 6 in the Appendix.

4. UNIVERSAL CONSTRUCTIBILITY

4.1. Information Hierarchies and Constructibility

These examples motivate the following issue. Partially order a collection of experiments so that higher experiments are Blackwell more informative than lower experiments. Is it possible to find signals that induce these experiments, and that satisfy the property that signals inducing higher experiments are refinements of signals inducing lower ones? We argued that this was in fact always possible when the ranking on experiments corresponded to the chain or the tree example above but was not necessarily possible when experiments were ranked as a diamond. The main result of this paper is an if-and-only-if characterization of the rankings of experiments such that suitable refinement-ordered signals exist for any experiments whose Blackwell order is consistent with the ranking.

We now provide the remaining definitions needed to state our result. An *information hierarchy* H is a finite partially ordered set (N, \geq) . We refer to elements of N as *nodes*. Given $n, n' \in N$, we say that n *covers* n' if $n > n'$ and there does not exist $n'' \in N$ with $n > n'' > n'$. The *graph of H* , denoted $G(H)$, is an undirected graph whose nodes are the elements of N with an edge between two nodes if one covers the other. $G(H)$ is a *forest* if there is at most one path between any two nodes.⁹

Fix an information hierarchy H and (Ω, μ_0) . An *experiment allocation* on H is a map that assigns an experiment to every node; experiment allocation β is *monotone* if $n \geq n'$ implies that $\beta(n) \succsim \beta(n')$, that is, higher nodes are Blackwell more informed than lower nodes. A *signal allocation* on H is a map that assigns a signal to every node; signal allocation σ is *monotone* if $n \geq n'$ implies that $\sigma(n) \supseteq \sigma(n')$, that is, signals associated with higher nodes refine signals associated with lower nodes. A signal allocation σ *induces* an experiment allocation β if, for all n , $\beta(n) = \langle \sigma(n) \rangle$. A (monotone) experiment allocation on H is *constructible* if there exists a monotone signal allocation that induces it.¹⁰

We say that H is *universally constructible* if, for every Ω and μ_0 , every monotone experiment allocation is constructible.

4.2. Main Result

We now present the main result of the paper:

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

To relate this result to our motivating examples, both the chain and the tree are examples of information hierarchies whose graphs are forests; hence, every monotone experiment allocation is constructible. The graph of the diamond, however, is not a forest

⁹It is standard to represent a partially ordered set as a directed graph, for example, as we do in Figure 1. Our main theorem relates constructibility of a hierarchy to a property of its undirected analogue; this is why we define $G(H)$ as the undirected version. Formally, the directed graph $\vec{G}(H)$ is the pair (N, \vec{E}) , where N is the set of nodes, $\vec{E} \subseteq N \times N$ is the set of directed edges, and $(n, n') \in \vec{E}$ if n covers n' . A directed path from n to n' is an alternating sequence of nodes and directed edges $(n_0, \vec{e}_0, \dots, n_{L-1}, \vec{e}_{L-1}, n_L)$, where $L > 0$, $n_0 = n$, $n_L = n'$, $n_l \in N$ for all $l \in \{0, \dots, L\}$, $\vec{e}_l = (n_l, n_{l+1}) \in \vec{E}$ for all $l \in \{0, \dots, L-1\}$, and $l \neq l' \Rightarrow \vec{e}_l \neq \vec{e}_{l'}$. $G(H)$ is the pair (N, E) , where N is the set of nodes, $E \subseteq \{e \subseteq N \mid |e| = 2\}$ is the set of undirected edges, and $\{n, n'\} \in E$ if n covers n' or n' covers n . An undirected path from n to n' is an alternating sequence of nodes and undirected edges $(n_0, e_0, \dots, n_{L-1}, e_{L-1}, n_L)$, where $L > 0$, $n_0 = n$, $n_L = n'$, $n_l \in N$ for all $l \in \{0, \dots, L\}$, $e_l = \{n_l, n_{l+1}\} \in E$ for all $l \in \{0, \dots, L-1\}$, and $l \neq l' \Rightarrow e_l \neq e_{l'}$.

¹⁰It is immediate that any constructible experiment allocation is monotone because $\pi \supseteq \pi' \Rightarrow \langle \pi \rangle \succsim \langle \pi' \rangle$.

because there are two paths from A to D , namely the path going through B and the path going through C ; hence, we were able to present a monotone experiment allocation that is not constructible.

A rigorous proof of Theorem 1 is in the [Appendix](#). The remainder of this section sketches the argument.

4.3. Proof Sketch: if

We begin by explaining why, if $G(H)$ is a forest, then any monotone experiment allocation can be constructed. We illustrate this for the special case when $G(H)$ is a tree. One can then obtain the general case by applying the same argument on each of the disjoint trees that make up the forest.

The proof relies on the following result, which is of independent information-theoretic interest. We say that π' is *statistically redundant* given $\hat{\pi}$ if $\langle \hat{\pi} \vee \pi' \rangle = \langle \hat{\pi} \rangle$, that is, observing $\hat{\pi}$ and π' yields the same beliefs as observing $\hat{\pi}$ only.¹¹

LEMMA 1: Fix signals $\bar{\pi}$ and π , with $\bar{\pi} \supseteq \pi$, and an experiment τ' that is Blackwell less informative than $\langle \pi \rangle$. There exists a signal π' that induces τ' and is statistically redundant given any signal that is between π and $\bar{\pi}$ in the refinement order.

To better understand the content of Lemma 1, it is helpful to note that the following, stronger, conjecture does not hold. One might think that, analogously to the aforementioned result of [Green and Stokey \(1978\)](#), for any π and τ' with τ' Blackwell less informative than $\langle \pi \rangle$, there exists a π' that induces τ' and is coarser than π . This is not the case.¹² Lemma 1 implies, however, that we can nonetheless find a π' that induces τ' and is statistically redundant given π , even though we cannot guarantee that $\pi \supseteq \pi'$. Moreover, the lemma further implies we can find a π' so that π' is statistically redundant given any $\hat{\pi}$ such that $\pi \subseteq \hat{\pi} \subseteq \bar{\pi}$.

Why do we need the upper bound $\bar{\pi}$? Once we fix any $\bar{\pi}$ and π , we can construct π' that is statistically redundant given $\bar{\pi}$, π , and anything in between. But, suppose we only fix π and consider an arbitrary π' . As long as π' is not coarser than π , there is always some refinement of π , say π^* , such that π' is not statistically redundant given π^* .

The proof of Lemma 1 is constructive. To see how the construction works, consider the example in [Figure 3](#). The state is either L or R . Each of the four rows represents a signal, that is, a partition of $\{L, R\} \times [0, 1]$. The two rectangles in each row represent the unit interval crossed with the two states. Each signal realization is indicated by its pattern. Note that $\bar{\pi} \supseteq \pi$, and while $\langle \pi \rangle \succ \langle \tilde{\pi} \rangle$, it is not the case that $\pi \supseteq \tilde{\pi}$ nor that $\langle \pi \vee \tilde{\pi} \rangle = \langle \pi \rangle$.

To establish the claim in Lemma 1, we need to construct a signal π' that induces the same beliefs as $\tilde{\pi}$, but is statistically redundant given π , and given any refinement of π up to $\bar{\pi}$. The bottom row illustrates such a construction. Each signal realization of π' corresponds to a signal realization of $\tilde{\pi}$, with the same likelihood in each state. However, the

¹¹If π' is statistically redundant given $\hat{\pi}$, we then also say $\hat{\pi}$ is *statistically sufficient* for π' . It may be helpful to contrast statistical sufficiency/redundancy with the stronger notion of refinement/coarsening. If $\hat{\pi} \supseteq \pi'$, we do not merely have that $\langle \hat{\pi} \vee \pi' \rangle = \langle \hat{\pi} \rangle$, we also have that $\langle \hat{\pi} \vee \pi' \vee \pi'' \rangle = \langle \hat{\pi} \vee \pi'' \rangle$ for any signal π'' . For example, in [Figure 3](#), π' is not a coarsening of π , but (as explained in the discussion below) π' is statistically redundant given π .

¹²For example, suppose that $\langle \pi \rangle \succ \tau'$ but the support of τ' has more elements than the number of signal realizations in π . Then, no π' that induces τ' could be a coarsening of π .

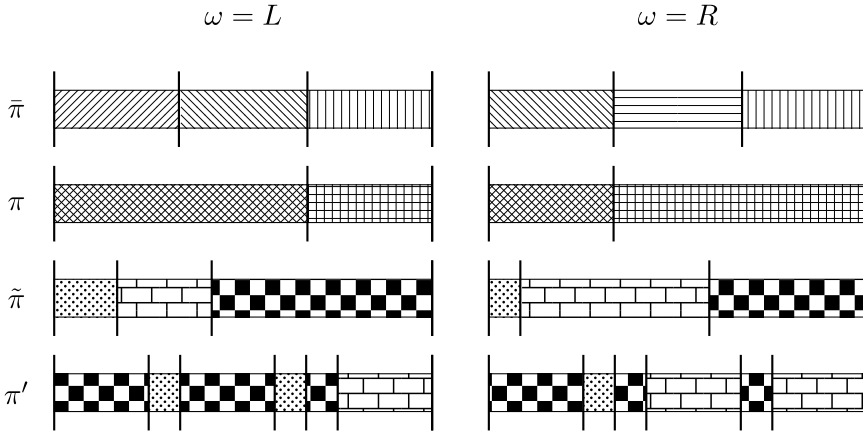


FIGURE 3.—Illustration of Lemma 1.

“locations” of the signal realizations in π' are rearranged so that the conditional probability of each signal realization of π' in state ω given $\bar{\pi}$ is (i) the same for $\omega = L$ and $\omega = R$, and (ii) the same for any elements of $\bar{\pi}$ that refine the same element of π . Property (i) ensures that π' is statistically redundant given $\bar{\pi}$, while (i) and (ii) together ensure that it is also redundant given any $\hat{\pi}$ s.t. $\bar{\pi} \supseteq \hat{\pi} \supseteq \pi$. The proof in the [Appendix](#) generalizes this construction.

With Lemma 1 in hand, we can now establish the universal constructibility of trees. Consider the tree hierarchy described in Section 3. Take some monotone experiment allocation β on this hierarchy. We construct a sequence of monotone signal allocations on progressively larger subsets of N , adding one node at a time. The signal assigned to a given node may be revised as the sequence progresses.

In Step 1, we pick an arbitrary node, say node D . Given $\beta(D)$, let $\sigma^1(D)$ be any signal that induces $\beta(D)$. In Step 2, we pick a node connected to D , node A . Node A is above node D and so we follow the same procedure as in the chain example to tentatively assign a suitable signal $\sigma^2(A)$ to A , with $\langle \sigma^2(A) \rangle = \beta(A)$ and $\sigma^2(A) \supseteq \sigma^1(D)$; the signal to D is unchanged in this step, that is, we let $\sigma^2(D) = \sigma^1(D)$. In Step 3, we pick an unassigned node connected to A or D , which here must be node B . Since B is below A , a complication in assigning a signal to B is that there may not exist a signal $\tilde{\pi}$ such that $\langle \tilde{\pi} \rangle = \beta(B)$ and yet $\sigma^2(A) \supseteq \tilde{\pi}$. This is where Lemma 1 comes into play. Applying Lemma 1 with $\pi = \bar{\pi} = \sigma^2(A)$ and $\tau = \beta(B)$, we conclude that there exists a signal π' that induces $\beta(B)$ and is statistically redundant with $\sigma^2(A)$, so that $\langle \sigma^2(A) \vee \pi' \rangle = \langle \sigma^2(A) \rangle$. We then set $\sigma^3(B) = \pi'$, we replace the initial assignment to A with $\sigma^3(A) = \sigma^2(A) \vee \pi'$, and we leave the signal at D unchanged at $\sigma^3(D) = \sigma^2(D)$. Finally, in Step 4, we move to the final node C , below the previously assigned node B . We now apply Lemma 1 with $\bar{\pi} = \sigma^3(A)$,¹³ $\pi = \sigma^3(B)$, and $\tau = \beta(C)$ to obtain a new signal π' that induces $\beta(C)$ and is statistically redundant with the previously assigned signals at nodes above C . We then set $\sigma^4(C) = \pi'$ and replace the signals of the previously assigned nodes above C as in Step 3.

This completes the construction of a monotone signal allocation inducing experiment allocation β . The full details of this procedure, applied to any hierarchy whose graph is a forest, are the heart of the formal proof of the if direction.

¹³In general, $\bar{\pi}$ is taken to be the join of all signals previously assigned to nodes above the new node, which in this case is simply $\sigma^3(A)$.

Which properties of a tree are used in this argument? When we apply this procedure to trees (or forests more generally), we can add the nodes in a way so that the new node n always covers or is covered by exactly one node n' to which we have previously assigned a signal. When n covers n' , we can apply Green and Stokey's (1978) result to obtain the new signal at n . When n is covered by n' , we can apply Lemma 1 instead. If the hierarchy is not a forest, however, no matter how we add the nodes, we will inevitably find ourselves in a situation where the new node covers or is covered by more than one node, in which case neither Green and Stokey's (1978) result nor Lemma 1 is of use.

4.4. Proof Sketch: Only if

We now sketch the proof of the other direction of Theorem 1: a hierarchy is universally constructible only if its graph is a forest. The proof consists of three main steps. Step 1 introduces a notion of a closed subhierarchy and establishes that a hierarchy is universally constructible only if its closed subhierarchies are as well. Step 2 shows that every hierarchy that is not a forest contains a closed subhierarchy that is one of two types. Step 3 shows that both of these types are not universally constructible, thus completing the proof.

Step 1: Given a hierarchy $H = (N, \geq)$, a subset of nodes $N' \subseteq N$ induces the hierarchy $H' = (N', \geq)$, which we refer to as a *subhierarchy* of H , with the partial order being the restriction of \geq to N' . Say that H' 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 2 in the Appendix shows that if H is universally constructible, then every closed subhierarchy of H is also universally constructible. The argument is as follows. Fix a closed subhierarchy H' . We show that any monotone experiment allocation β on H' can be extended to a monotone experiment allocation on all of H . This step uses the hypothesis that H' is closed, since it means that in extending β , we never have to fill in beliefs for "in between" nodes in H' . As H is universally constructible, there is a monotone signal allocation on H that induces the extension of β , and the restriction of this signal allocation to N' is a monotone signal allocation on H' that induces β . Since β is arbitrary, H' is universally constructible.

Step 2: Say that a hierarchy H is *cyclic* if $G(H)$ is not a forest. Suppose H is cyclic. Given that N is finite and H is cyclic, it follows that H contains a *minimal cyclic closed subhierarchy (MCC)*, that is, a subhierarchy $H' = (N', \geq)$ such that: H' is cyclic, closed, and there does not exist $N'' \subsetneq N'$ such that $H'' = (N'', \geq)$ is cyclic and closed. Any MCC H' must fall into one of the following categories:

- H' is a *union of non-comparable paths (UNP)*: It contains a maximal node \bar{n} and a minimal node \underline{n} , and its graph consists of at least two paths between \bar{n} and \underline{n} . Moreover, nodes in $N \setminus \{\bar{n}, \underline{n}\}$ are comparable only if they are in the same path. The diamond is an example of a UNP; another example is depicted in Figure 4(a).
- H' is a *crown*: It contains exactly four nodes, which we label F (ather), M (other), S (on), D (aughter), and the partial order consists of $F \geq S$, $F \geq D$, $M \geq S$, and $M \geq D$. A crown is depicted in Figure 6(a).

This taxonomy takes considerable effort to prove formally, but the high level argument is as follows. An MCC H' contains maximal nodes, which are not covered by any other node, and minimal nodes, which do not cover any other nodes. We distinguish two cases on these maximal and minimal nodes.

In the first case, every maximal node in H' covers every minimal node. Then, the fact that H' is minimal, together with the existence of a cycle, implies that H' contains exactly four nodes and is in fact a crown.




FIGURE 4.—For adjacent nodes, the node that is higher in the figure covers the node that is lower. (a) A UNP. Nodes A , B , and C are in different paths, and therefore are not comparable. (b) A between set which is not an MCC. There are two overlapping (directed) paths that go from \bar{n} to \underline{n} , one that goes through B and one that does not. There is a smaller cyclic closed subhierarchy, namely the nodes between A and \underline{n} .

Alternatively, there is a maximal node \bar{n} that does not cover a minimal node \underline{n} . Then, we show that H' is a UNP. To see this, note the following two subcases. First, it may be that H' is simply the set of nodes that are between \bar{n} and \underline{n} , that is, a *between set*. Then, H' consists of a series of directed paths between \bar{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 4(b), violating the fact that H' is minimal. Thus, nodes must not be comparable across paths and H' is a UNP. The second subcase is that H' is not a between set. Lemma 4 in the Appendix shows that then every cycle in H' must contain every node in N' . 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' , which would contradict Lemma 4. As a result, H' is a UNP.

Step 3: Finally, we show that UNPs and crowns are not universally constructible. The former is demonstrated in Lemma 6 in the Appendix using a similar construction as we used for the diamond in Section 3. Indeed, the diamond is an example of a UNP in which there are exactly two undirected paths from \bar{n} to \underline{n} .

For the crown, we construct a particular monotone experiment allocation that is not constructible: Let $\Omega = \{0, 1, 2\}$ and $\mu_0(\omega = 0) = \frac{1}{2}$, $\mu_0(\omega = 1) = \frac{1}{4}$, and $\mu_0(\omega = 2) = \frac{1}{4}$. We represent each belief as a pair (x, y) in the triangle $\{(x, y) | x \geq 0, y \geq 0, x + y \leq 1\}$, where $\Pr(\omega = 1) = x$ and $\Pr(\omega = 2) = y$. Consider the experiment allocation β that assigns to F , M , S , and D the experiments indicated in Figure 5(a). Each belief realization is a circle, the letters inside the circle indicate the nodes for which this belief is in the support, and the area of the circle is the likelihood of the belief (which is the same for all nodes attached to the belief). For example, the belief $(1/4, 1/4)$ can be realized for F , S , and D , with likelihood of $1/2$ in each case. It is easy to see from Figure 5 that β assigns higher nodes more informative signals. For example, going from D to M (Figure 5(b)), the mass on the east and west beliefs stays the same, but the mass on the central belief is spread out in the north-south directions.




FIGURE 5.—The crown is not universally constructible.

Lemma 7 in the Appendix formally defines this experiment allocation and shows that it is not constructible. To see what goes wrong, suppose there is a monotone signal allocation that induces β . Such a signal allocation induces a joint distribution over beliefs at all four nodes. Consider first the joint distribution of beliefs between D , M , and S . The center D belief occurs if and only if the M belief is north or south (Figure 5(b)). The center belief of S occurs if and only if the M belief is east or west (Figure 5(c)). Hence, the probability of both D and S having the center belief is zero. Now consider the joint distribution of beliefs between D , F , and S . The belief at D is in the center if and only if the belief at F is in the center (Figure 5(d)), and the belief at S is in the center if and only if the belief at F is in the center (Figure 5(e)). So, the probability that both D and S have the center belief is positive, yielding a contradiction.

5. DISCUSSION

5.1. Binary States

A key step in our proof of the only-if direction of Theorem 1 shows that the crown is not universally constructible. Specifically, Figure 5 presents a particular experiment allocation involving three states that is not constructible.

It is natural to ask whether there is an example of a non-constructible allocation on the crown using two states. No such example exists: when the state space is binary, every monotone experiment allocation on the crown is constructible. The reason is as follows. Take an arbitrary monotone experiment allocation β on $\{F, M, S, D\}$. It is known that the set of experiments under the Blackwell order is a lattice when the state space is binary (Kertz and Rösler (2000), Müller and Scarsini (2006)). Thus, there exists an experiment $\beta(S) \vee \beta(D)$ such that $\beta(S) \vee \beta(D) \succeq \beta(S)$, $\beta(S) \vee \beta(D) \succeq \beta(D)$, and $\tau \succeq \beta(S) \vee \beta(D)$ for any $\tau \succeq \beta(S), \beta(D)$. We can therefore expand the crown hierarchy by adding a T (utor) whom F (ather) and M (other) hired to oversee the children: $F \geq T$, $M \geq T$, $T \geq S$, and $T \geq D$. We refer to this hierarchy as the *cross*, which is depicted




FIGURE 6.—The crown and cross hierarchies.

in Figure 6(b). Let $\hat{\beta}$ be the experiment allocation on the cross that sets $\hat{\beta}(n) = \beta(n)$ for $n \in \{F, M, S, D\}$ and $\hat{\beta}(T) = \beta(S) \vee \beta(D)$. The allocation $\hat{\beta}$ is monotone by construction. The cross is a forest, and hence by Theorem 1, it is universally constructible. Thus, there is a monotone signal allocation $\hat{\sigma}$ that induces $\hat{\beta}$. Restricting $\hat{\sigma}$ to $\{F, M, S, D\}$ yields a monotone signal allocation on the crown that induces β .

This example illustrates that the binary state space is special. With three or more states, every monotone experiment allocation is constructible if and only if the hierarchy is a forest. With binary states, being a forest is sufficient for constructibility, but it is not necessary. A natural direction for future research is to characterize necessary and sufficient conditions for every monotone experiment allocation to be constructible for the case of binary states.

5.2. Beyond the Refinement Order

As we mentioned at the outset, there are various things one might mean by “ A is more informed than B .” One is that A ’s signal is Blackwell more informative than B ’s, that is, $\langle \pi_A \rangle \succeq \langle \pi_B \rangle$. Another is that A has observed all of B ’s information, that is, $\pi_A \supseteq \pi_B$. These, however, are not the only economically relevant comparisons of signals. For instance, it might be that A knows B ’s belief about the state. Or, it might be that if A observed B ’s information, A ’s belief about the state of the world would not change, that is, that π_B is statistically redundant given π_A . In a companion paper (Brooks, Frankel, and Kamenica (2022)), we explore these and other relations on signals in more detail. In this section, we explain how our results extend to a large class of notions of “more informed,” including those just mentioned.

First, consider the belief-martingale relation on Π , denoted \mathcal{B} , defined by $\pi \mathcal{B} \pi'$ if $\mathbb{E}[\tilde{\mu}_\pi | \tilde{\mu}_{\pi'}] = \tilde{\mu}_{\pi'}$ for any prior. (We use the word relation rather than order because \mathcal{B} is not transitive.)¹⁴ Note that the refinement order implies the belief-martingale relation, which in turn implies the Blackwell order:

$$\pi \supseteq \pi' \quad \Rightarrow \quad \pi \mathcal{B} \pi' \quad \Rightarrow \quad \langle \pi \rangle \succeq \langle \pi' \rangle.$$

Many economically relevant comparisons of signals are “in between” refinement and the belief-martingale relation, that is, they are weaker than the former and stronger than

¹⁴In Brooks, Frankel, and Kamenica (2022), we illustrate that \mathcal{B} is not transitive and that, given two signals π and π' , the equality $\mathbb{E}[\tilde{\mu}_\pi | \tilde{\mu}_{\pi'}] = \tilde{\mu}_{\pi'}$ might hold for one prior but not for another.

the latter. For example, the two aforementioned comparisons— A knowing B 's belief or B 's signal being statistically redundant given A 's—fall in this in-between category. Formally, we say a binary relation \mathcal{R} on Π is *proper* if $\pi \supseteq \pi' \Rightarrow \pi \mathcal{R} \pi' \Rightarrow \pi \mathcal{B} \pi'$.

Our main result can be extended by replacing the refinement order with any proper relation. In particular, say a signal allocation σ on H is \mathcal{R} -*monotone* if $n \geq n'$ implies that $\sigma(n) \mathcal{R} \sigma(n')$. The experiment allocation β on H is \mathcal{R} -*constructible* if β is induced by some \mathcal{R} -monotone signal allocation on H . A hierarchy is said to be \mathcal{R} -*universally constructible* if every monotone experiment allocation is \mathcal{R} -constructible for any Ω and μ_0 . With this terminology, we have the following:

THEOREM 2: *Fix any proper relation \mathcal{R} . A hierarchy H is \mathcal{R} -universally constructible if and only if its graph is a forest.*

The if direction follows immediately from Theorem 1, since \supseteq -universal constructibility implies \mathcal{R} -universal constructibility for any proper \mathcal{R} . For the only if direction, the key is to recognize that the proofs of Lemmas 6 and 7 in the Appendix establish not only that UNPs and crowns are not \supseteq -universally constructible, but also that they are not \mathcal{B} -universally constructible. With that change, the remainder of the proof of the only if direction of Theorem 1 (replacing every instance of “universally constructible” with “ \mathcal{B} -universally constructible”) establishes that if $G(H)$ is not a forest, it is not \mathcal{B} -universally constructible. A fortiori, it is not \mathcal{R} -universally constructible for any proper \mathcal{R} .

6. APPLICATIONS

6.1. Rationalizing Reaction to Unknown Sources of Information

Consider an 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 a decision maker has access to a set of information sources $\{x_1, x_2, \dots, x_M\}$. Suppose further that our data set $\mathcal{D} = \{\tau_S\}_{S \subseteq \{x_1, \dots, x_M\}}$ tells us the distribution of beliefs conditional on observing any non-empty subset of information sources. When can we rationalize a given data set \mathcal{D} in the sense that we can associate each information source x_i with some signal (i.e., an element of Π) and conclude that belief formation is consistent with Bayes's rule?

To be rationalized, belief distributions (i.e., experiments) in \mathcal{D} have to satisfy two obvious properties. First, there is *Bayes plausibility*: the average belief cannot differ across sets of information sources, that is, $\mathbb{E}_{\tau_S}[\mu] = \mathbb{E}_{\tau_{S'}}[\mu]$ for any two subsets S and S' . Second, there is *Blackwell monotonicity*: observing a superset of sources necessarily induces a more dispersed distribution of beliefs, that is, 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, when there are three or more sources, the answer is no.¹⁵ In this case, 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 sources $S \subseteq \{x_1, \dots, x_M\}$ 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 7, the graph

¹⁵When there are only two sources, $M = 2$, 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.




FIGURE 7.—The set-inclusion hierarchy.

of this information hierarchy is not a forest. By Theorem 1, this means that there is some monotone experiment allocation on H , call it β , that cannot be induced by any monotone signal allocation on H . Now, we can associate with this β a data set $\mathcal{D} = \{\tau_S\}_{S \subseteq \{x_1, \dots, x_M\}}$ by setting $\tau_S = \beta(n_S)$. 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 be monotone (since $S' \subseteq S$ implies $\bigvee_{x_i \in S'} \pi(x_i) \subseteq \bigvee_{x_i \in S} \pi(x_i)$), contradicting Theorem 1. Thus, we know that there are data sets that satisfy Bayes plausibility and Blackwell monotonicity, yet cannot be rationalized.

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

6.2. Information Design

Suppose a sender provides information to a set of agents. Moreover, she faces certain types of monotonicity constraints such as some agents must know the beliefs of some other agents or some agents must have access to others' information. Then, we can think of agents as elements of an information hierarchy, and the information design problem consists of selecting a (suitably monotone) signal allocation on this hierarchy. Our results shed light on how such monotonicity constraints affect the information design problem.

For example, consider organizations. A long literature in organization economics emphasizes the importance of the hierarchical structure of managerial relationships (Williamson (1967)). One important aspect of organizational design is deciding how much information to provide—about individuals' prospects for promotion, about the overall performance of the organization, 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 the information that is available to her subordinates, that is, that the allocation of signals within the organization should be monotone with respect to the management structure, so that the signal of a superior refines the signals of their subordinates. This constraint interacts with the organization structure. Theorem 1 implies that, if an organization has the feature that every subordinate has at most one superior, 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. Furthermore,

if the information designer's objectives only depend on each agent's experiment, then one could reformulate the information design problem in terms of the choice of monotone experiment allocation, rather than choosing monotone signal allocations directly. With richer managerial relationships,¹⁶ however, Theorem 1 also tells us there could be desirable allocations of information which are incompatible with the monotonicity constraint, even though they provide (Blackwell) more information to those higher up in the organization.

7. RELATED LITERATURE

We now discuss the connection between our results and related literature in probability theory and in information economics.¹⁷

7.1. Connections to Probability Theory

Mathematically, our results are close to a literature in probability theory on the relationship between orders on probability distributions and stochastic processes. Because the connection to this literature is technical, our discussion is somewhat formal and detailed.

Fix two distributions $\tau, \tau' \in \Delta(\mathbb{R}^k)$. We say that τ is greater than τ' in the convex order if, for every convex function $\phi: \mathbb{R}^k \rightarrow \mathbb{R}$, $\int \phi d\tau \geq \int \phi d\tau'$. The Sherman–Stein theorem says that the following are equivalent: (i) τ is greater than τ' in the convex order; (ii) τ can be obtained from τ' via mean-preserving spreads. Blackwell (1953, Section 3) observed that his result is implied by the Sherman–Stein theorem, since we can interpret τ as an experiment and the convex function ϕ as the maximum utility obtained in a decision problem given a belief. Importantly, the mean-preserving spread from τ to τ' is a probability transition kernel which defines a martingale on \mathbb{R}^k . Thus, an equivalent way of stating (ii) is that there exists a two-period martingale on \mathbb{R}^k such that the marginals are τ and τ' .

This equivalence was later generalized by Strassen (1965, Theorem 8): the sequence $\tau_1, \dots, \tau_M \in \Delta(\mathbb{R}^k)$ is increasing in the convex order if and only if there exists an M -period martingale with marginals τ_1, \dots, τ_M . This result is closely related to Green and Stokey (1978) and, more broadly, whether experiments can be represented as partitions: When the set of realizations of the martingale is finite, the martingale generates a filtration $(\Sigma_m)_{m=1}^M$, where each Σ_m is generated by a finite partition π_m of the underlying probability space, and the partitions are increasing in the refinement order. We can therefore view the probability space itself as an expanded state space, so that the π_m are refinement-ordered signals that induce the given experiments. In our terminology, Strassen's theorem implies that the chain is universally constructible.

There is a substantial literature that further generalizes and strengthens Strassen's theorem. See Hirsch, Profeta, Royonette, and Yor (2011) for a survey.¹⁸ To our knowledge, Juillet (2016) is the only paper in this literature to study collections of distributions that are indexed on partially ordered sets. He presented examples that illustrate various ways in which Strassen's theorem does not generalize. Most closely related to the present study,

¹⁶For instance, suppose that the CEO oversees two middle managers who share the oversight of an employee.

¹⁷Our inquiry also leads us to answer a pure graph-theoretic question regarding whether a partially ordered set contains subsets of a particular form. Similar questions have been studied in combinatorics and graph theory (e.g., Lu (2014)); within economics, results of this form are used by Curello and Sinander (2019) to study rankings on preference relations.

¹⁸We thank Ludvig Sinander and Aniko Oery for drawing our attention to this literature.

Juillet (2016, Section 4.1) constructed a collection $(\tau_x)_{x \in X}$, where X is the diamond and $\tau_x \in \Delta(\mathbb{R})$, such that the collection is increasing in the convex order but there does not exist a real-valued martingale on X with the given marginals. The distributions in Juillet’s example are similar in spirit to (but distinct from) our example of a non-constructible belief allocation given in Section 3.¹⁹ Our Theorem 1 implies that the equivalence between $(\tau_x)_{x \in X}$ increasing in the convex order and induced by a martingale holds if and only if X is a forest.

Also related is Fill, Allen, and Machida (2001), who studied an analogue of the problem of Juillet (2016), but where the convex order is replaced by a generalization of first-order stochastic dominance. Specifically, Fill, Allen, and Machida (2001) considered collections of distributions indexed on a partially ordered set $(\tau_x)_{x \in X}$, where each τ_x is a distribution on a different partially ordered set S . They defined a notion of stochastic monotonicity for distributions on S that reduces to first-order stochastic dominance when S is totally ordered. A pair (X, S) is *monotonicity equivalent* if the following condition holds: a collection $(\tau_x)_{x \in X}$ is stochastically monotone if and only if there is a distribution over non-decreasing functions $f : X \rightarrow S$ such that τ_x is the marginal distribution of $f(x)$. Monotonicity equivalence is analogous to universal constructibility, but with stochastic monotonicity in place of the Blackwell order. In the special case that S contains a subposet that is either a cycle or a crown, Fill, Allen, and Machida (2001) showed that (X, S) is monotonicity equivalent if and only if X is a forest. The cases where X is a diamond or a crown also feature prominently in their proof, which follows a similar strategy of showing that if X is not a forest, then it contains a sub-poset on which monotonicity equivalence fails, and then extending the counterexample to all of X .

7.2. Connections to Information Economics

Our paper connects to various threads in the economics literature studying information and beliefs. In addition to the aforementioned classical literature on information orders, Bergemann and Morris (2016) studied an extension of the Blackwell order to type spaces, and Mu, Pomatto, Strack, and Tamuz (2019) considered comparisons of repeated experiments.

We contribute to the literature on higher order beliefs (Harsanyi (1967), Mertens and Zamir (1985), Brandenburger and Dekel (1993)). Specifically, we characterize when restrictions of the form “player i knows player j ’s type” place constraints on i ’s and j ’s first-order beliefs, beyond the obvious constraint that i must be Blackwell more informed than j .²⁰

We also study how signals can be combined to produce more informative signals. Gentzkow and Kamenica (2017) studied this issue in the context of a communication game with a receiver who combines information provided by multiple senders. Börgers, Hernando-Veciana, and Krämer (2013) studied the interaction between signals from the perspective of whether signals are substitutes or complements.

We also contribute to the growing literature on information design (Kamenica and Gentzkow (2011), Bergemann and Morris (2016)). Arieli, Babichenko, Sandomirskiy, and

¹⁹It is straightforward to adapt Juillet’s example to show that the diamond is not universally constructible. But unlike our example from Section 3, his example is not easily adapted for general UNPs.

²⁰In the language of type spaces, “player i knows player j ’s type” is analogous to the refinement order. As discussed in Section 5.2, we also study relations that are weaker than refinement.

Tamuz (2021) characterized feasible joint belief distributions of a group of agents in a binary state case. Mathevet and Taneva (2020) analyzed the implications of information design for organizational structure.

Finally, Shmaya and Yariv (2016) studied behavior by subjects in lab experiments when the experimental protocol potentially generates information. Their analysis led them to a question that is related to the result we present in Section 6.1. In particular, given data on how an agent behaves following any sequence of signal realizations, Shmaya and Yariv derived necessary and sufficient conditions for the behavior to be consistent with Bayesian updating if the data generating process, the agent’s preferences, and the agent’s prior are not known. On a similar topic, Molavi (2021) studied which belief dynamics are consistent with Bayesian updating, but allowed for the possibility that the agent has an incorrect model of the data generating process.

8. CONCLUSION

We conclude with a brief discussion of an important direction for future work. Our analysis has focused on the question of when an information hierarchy is universally constructible. That condition requires that every monotone experiment allocation is constructible. Of course, even when not every monotone experiment allocation is constructible, some are. A natural goal would therefore be to characterize, for a given information hierarchy, the set of constructible monotone experiment allocations. Such a characterization would be of particular interest in the case of the set-inclusion hierarchy (Figure 7), which would identify the exact sufficient and necessary conditions for rationalizability of reactions to unknown sources of information, as discussed in Section 6.1.

APPENDIX: PROOFS

A.1. Proof of Lemma 1

PROOF OF LEMMA 1: Let $\tilde{\pi}$ be a signal s.t. $\langle \tilde{\pi} \rangle = \tau'$. Since $\langle \pi \rangle \succsim \tau'$, there exists a garbling $g : \pi \times \tilde{\pi} \rightarrow [0, 1]$ such that $\sum_{\tilde{s} \in \tilde{\pi}} g(s, \tilde{s}) = 1 \ \forall s \in \pi$, and $p(\tilde{s}|\omega) = \sum_{s \in \pi} g(s, \tilde{s})p(s|\omega)$. For every $\bar{s} \in \tilde{\pi}$, let $\underline{s}(\bar{s})$ denote the element of π s.t. $\bar{s} \subseteq \underline{s}(\bar{s})$. (This element exists since $\tilde{\pi} \supseteq \pi$.) Now, $\forall \bar{s} \in \tilde{\pi}$, let $\{X_{\tilde{s}}^{\bar{s}}\}_{\tilde{s} \in M^{\bar{s}}}$ be a partition of \bar{s} s.t. $\forall \omega, \lambda(\{x|(\omega, x) \in X_{\tilde{s}}^{\bar{s}}\}) = \lambda(\{x|(\omega, x) \in \bar{s}\})g(\underline{s}(\bar{s}), \tilde{s})$, where $M^{\bar{s}} = \{\tilde{s} \in \tilde{\pi} | g(\underline{s}(\bar{s}), \tilde{s}) > 0\}$. Such a partition exists because $\sum_{\tilde{s} \in \tilde{\pi}} g(\underline{s}(\bar{s}), \tilde{s}) = 1$ for all $\underline{s}(\bar{s}) \in \pi$. Let $\pi' = \{Z^{\tilde{s}}\}_{\tilde{s} \in \tilde{\pi}}$ with $Z^{\tilde{s}} = \bigcup_{\bar{s} \in \tilde{\pi} \text{ s.t. } \tilde{s} \in M^{\bar{s}}} X_{\tilde{s}}^{\bar{s}}$. We now show that π' satisfies (i) and (ii). To show (i), it suffices to show that $p(Z^{\tilde{s}}|\omega) = p(\tilde{s}|\omega)$ for every \tilde{s} and ω . We have

$$\begin{aligned} p(Z^{\tilde{s}}|\omega) &= \lambda\left(\left\{x|(\omega, x) \in \bigcup_{\substack{\bar{s} \in \tilde{\pi} \\ \text{s.t. } \tilde{s} \in M^{\bar{s}}}} X_{\tilde{s}}^{\bar{s}}\right\}\right) \\ &= \sum_{\substack{\bar{s} \in \tilde{\pi} \\ \text{s.t. } \tilde{s} \in M^{\bar{s}}}} \lambda(\{x|(\omega, x) \in X_{\tilde{s}}^{\bar{s}}\}) \\ &= \sum_{\substack{\bar{s} \in \tilde{\pi} \\ \text{s.t. } \tilde{s} \in M^{\bar{s}}}} \lambda(\{x|(\omega, x) \in \bar{s}\})g(\underline{s}(\bar{s}), \tilde{s}) \\ &= \sum_{\bar{s} \in \tilde{\pi}} \lambda(\{x|(\omega, x) \in \bar{s}\})g(\underline{s}(\bar{s}), \tilde{s}) \end{aligned}$$

$$\begin{aligned}
 &= \sum_{s \in \pi} \sum_{\bar{s} \text{ s.t. } \underline{g}(\bar{s})=s} \lambda(\{x | (\omega, x) \in \bar{s}\}) g(\underline{g}(\bar{s}), \hat{s}) \\
 &= \sum_{s \in \pi} g(s, \hat{s}) \sum_{\bar{s} \text{ s.t. } \underline{g}(\bar{s})=s} \lambda(\{x | (\omega, x) \in \bar{s}\}) \\
 &= \sum_{s \in \pi} g(s, \bar{s}) \lambda(\{x | (\omega, x) \in s\}) \\
 &= \sum_{s \in \pi} g(s, \bar{s}) p(s | \omega) \\
 &= p(\bar{s} | \omega).
 \end{aligned}$$

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

$$\frac{p(s_i^\vee | \omega)}{p(s_i^\vee | \omega')} = \frac{p(\hat{s} | \omega)}{p(\hat{s} | \omega')}.$$

Consider some s_i^\vee . Note that there exists $\underline{s} \in \pi$ with $\hat{s} \subseteq \underline{s}$ since $\hat{\pi} \supseteq \pi$. Let $Q = \{\bar{s} \in \bar{\pi} | \bar{s} \subseteq \hat{s}\}$. Since $\bar{\pi} \supseteq \hat{\pi}$, for every ω , $\lambda(x | (\omega, x) \in \hat{s}) = \sum_{\bar{s} \in Q} \lambda(x | (\omega, x) \in \bar{s})$. Note that $\bar{s} \subseteq \underline{s}$ for all $\bar{s} \in Q$. Now, we know that $s_i^\vee = s' \cap \hat{s}$ for some $s' \in \pi'$. By definition of π' , we know that $s' = \bigcup_{\bar{s} \in \bar{\pi} \text{ s.t. } \bar{s} \in M^{\bar{s}}} X_{\bar{s}}^{\bar{s}}$ for some $\bar{s} \in \bar{\pi}$. Hence,

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

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

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

$$\begin{aligned}
 &= g(\underline{s}, \tilde{s}) \sum_{\bar{s} \in Q} \lambda(\{x | (\omega, x) \in \bar{s}\}) \\
 &= g(\underline{s}, \tilde{s}) \lambda(\{x | (\omega, x) \in \hat{s}\}) \\
 &= g(\underline{s}, \tilde{s}) p(\hat{s} | \omega).
 \end{aligned}$$

Hence,

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

which completes the proof of Lemma 1.

Q.E.D.

A.2. Proof of Theorem 1: if

Let H be an information hierarchy and suppose $G(H)$ is a forest. Let β be a monotone experiment allocation on H . We will construct a monotone signal allocation that induces β . To do so, we construct a sequence of subhierarchies of H , adding nodes of H one by one, until we reach the full hierarchy H . At each step, we assign a signal to the newly added node and potentially reassign the signals allocated to the previously added nodes.

We begin with some notation and terminology. A *construction procedure* f is a bijection from $\{1, \dots, |N|\}$ to N that specifies the order in which the nodes are added. Let $N_l^f = \{f(1), \dots, f(l)\}$. 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

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

Now, consider a construction procedure f of the following form. Let $f(1)$ be any node in N . For $l \in \{2, 3, \dots, |N|\}$, let $f(l)$ be an arbitrary element of $CoveredBy(N_{l-1}^f) \cup Covering(N_{l-1}^f) \cup Disconnected(N_{l-1}^f)$. 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 $G(H)$ between $f(l)$ and nodes in N_{l-1}^f .

PROOF OF CLAIM 1: Suppose toward contradiction that $f(l)$ has an edge in $G(H)$ with distinct $n, n' \in N_{l-1}^f$. Since n and n' both have an edge with $f(l)$, they must belong to the same tree in $G(H)$. Moreover, there must be a path between n and n' in $G((N_{l-1}^f, \geq))$. 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 $f(l') \in CoveredBy(N_{l'-1}^f) \cup Covering(N_{l'-1}^f)$, which in turn means that there is a path from $f(l')$ to \underline{n} in $G((N_{l'-1}^f, \geq))$ and thus in $G((N_{l-1}^f, \geq))$. Hence, there is a path from both n and n' to \underline{n} and thus a path between n and n' in $G((N_{l-1}^f, \geq))$. So, there must be a path between n and n' in $G(H)$ 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, $G(H)$ is a forest, so there cannot be multiple paths between two nodes; we have reached a contradiction. *Q.E.D.*

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 $\sigma^l : N_l^f \rightarrow \Pi$. We proceed by induction and show that, as long as the signals previously allocated to nodes in N_{l-1}^f are monotone on (N_{l-1}^f, \geq) and induce appropriate experiments (i.e., for all $m \in N_{l-1}^f$, $\langle \sigma^{l-1}(m) \rangle = \beta(m)$), then the σ^l we specify is monotone on (N_l^f, \geq) and induces appropriate experiments.

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 experiments. For $l \geq 2$, there are three cases: $f(l) \in \text{CoveredBy}(N_{l-1}^f)$, $f(l) \in \text{Covering}(N_{l-1}^f)$, and $f(l) \in \text{Disconnected}(N_{l-1}^f)$.

We first consider the case $f(l) \in \text{Covering}(N_{l-1}^f)$. 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 $\beta(f(l)) \succeq \beta(\bar{m})$, there exists some $\pi \supseteq \sigma^{l-1}(\bar{m})$ such that $\langle \pi \rangle = \beta(f(l))$ (which follows from Theorem 1 of Green and Stokey (1978)). We set $\sigma^l(f(l)) = \pi$ and we keep the signal allocation for nodes in N_{l-1}^f unchanged, that is, $\sigma^l(m) = \sigma^{l-1}(m)$ for all $m \in N_{l-1}^f$. It is clear that σ^l induces appropriate experiments (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 is monotone on N_l^f . Consider any $m, m' \in N_l^f$ such that $m > m'$. Since $f(l) \in \text{Covering}(N_{l-1}^f)$, either $m, m' \in N_{l-1}^f$ or $f(l) = m$. In the former case, we know $\sigma^l(m) = \sigma^{l-1}(m) \supseteq \sigma^{l-1}(m') = \sigma^l(m')$ by the inductive hypothesis. If $f(l) = m$, we know $f(l) > \bar{m} \geq m'$. By the inductive hypothesis, $\sigma^l(\bar{m}) = \sigma^{l-1} \supseteq \sigma^{l-1}(m') = \sigma^l(m')$ and thus $\sigma^l(f(l)) \supseteq \sigma^l(\bar{m}) \supseteq \sigma^l(m')$. That completes the proof for this case.

Now consider the case where $f(l) \in \text{CoveredBy}(N_{l-1}^f)$. Let \underline{m} be the node in N_{l-1}^f that covers $f(l)$. Denote $\tau' = \beta(f(l))$, $\pi = \sigma^{l-1}(\underline{m})$, and $\bar{\pi} = \bigvee_{m \in N_{l-1}^f} \sigma^{l-1}(m)$. By Lemma 1, we know $\exists \pi'$ such that (i) $\langle \pi' \rangle = \tau'$, and (ii) $\forall \hat{\pi}$ s.t. $\bar{\pi} \supseteq \hat{\pi} \supseteq \pi$, $\langle \hat{\pi} \vee \pi' \rangle = \langle \hat{\pi} \rangle$. 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 is monotone on (N_l^f, \geq) and induces appropriate experiments. 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, $\sigma^{l-1}(m) \supseteq \sigma^{l-1}(\underline{m}) = \pi$; moreover, $\bar{\pi} = \bigvee_{m' \in N_{l-1}^f} \sigma^{l-1}(m') \supseteq \sigma^{l-1}(m)$; hence, $\langle \sigma^l(m) \rangle = \langle \sigma^{l-1}(m) \vee \pi' \rangle = \langle \sigma^{l-1}(m) \rangle$.

For $m \in N_{l-1}^f$ s.t. $m \not\geq f(l)$, $\langle \sigma^l(m) \rangle = \langle \sigma^{l-1}(m) \rangle$. 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 $m \in N_{l-1}^f$. We now need to show that σ^l is monotone. Consider any $m, m' \in N_l^f$ 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) $\sigma^{l-1}(m) \supseteq \sigma^{l-1}(m')$, we know that $\sigma^{l-1}(m) \vee \pi' \supseteq \sigma^{l-1}(m') \vee \pi'$, and hence $\sigma^l(m) \supseteq \sigma^l(m')$. The second case is where $m \geq f(l)$ and $m' \not\geq f(l)$. Then, $\sigma^l(m) = \sigma^{l-1}(m) \vee \pi'$ and $\sigma^l(m') = \sigma^{l-1}(m')$. Since (by the inductive hypothesis) $\sigma^{l-1}(m) \supseteq \sigma^{l-1}(m')$, we have that $\sigma^l(m) = \sigma^{l-1}(m) \vee \pi' \supseteq \sigma^{l-1}(m) \supseteq \sigma^{l-1}(m') = \sigma^l(m')$. Finally, suppose that $m \not\geq f(l)$ and $m' \not\geq f(l)$. Then, $\sigma^l(m) = \sigma^{l-1}(m)$ and $\sigma^l(m') = \sigma^{l-1}(m')$. Since (by the inductive hypothesis) $\sigma^{l-1}(m) \supseteq \sigma^{l-1}(m')$, we have that $\sigma^l(m) \supseteq \sigma^l(m')$.

Finally, suppose $f(l) \in \text{Disconnected}(N_{l-1}^f)$. 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_{l-1}^f

unchanged, that is, $\sigma^l(m) = \sigma^{l-1}(m)$ for all $m \in N_{l-1}^f$. It is clear that σ^l induces appropriate experiments (by the inductive hypothesis for $m \in N_{l-1}^f$ and by construction for $f(l)$). Since $f(l)$ is not comparable to any node in N_{l-1}^f , the fact that σ^{l-1} is a signal allocation on (N_{l-1}^f, \geq) implies that σ^l is monotone on (N_l^f, \geq) . This completes the proof. \square

A.3. Proof of Theorem 1: Only if

LEMMA 2: *If an information hierarchy H is universally constructible and H' is a closed subhierarchy of H , then H' is universally constructible.*

PROOF OF LEMMA 2: Suppose H is an information hierarchy that is universally constructible under \mathcal{B} . Fix a state space Ω and a prior μ_0 . Suppose $H' = (N', \geq)$ is a closed subhierarchy of H . Let β' be a monotone experiment allocation on H' . We need to construct a monotone signal allocation on H' that induces β' . Let $\beta : N \rightarrow \Delta(\Delta(\Omega))$ defined 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 $\beta(n) = \bar{\tau}$; 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 show that $\beta(n) \succsim \beta(n')$ by considering four exhaustive cases:

If n and n' are both in N' , this follows from the fact that β' is monotone on H' .

If n and n' are both not in N' , consider two subcases. If $\exists n'' \in N'$ such that $n > n''$, then $\beta(n) = \bar{\tau} \succsim \beta(n')$. 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 $n' \in N'$, $\beta(n) = \bar{\tau} \succsim \beta(n')$.

Finally, if $n \in N'$ and $n' \notin N'$, then there cannot exist an $n'' \in N'$ with $n' > n''$. If such an n'' did exist, then since H' is closed and $n, n'' \in N'$, we would have that $n' \in N'$, a contradiction. Thus, $\beta(n) \succsim \beta(n') = \underline{\tau}$. Q.E.D.

Since β is a monotone experiment allocation on H and H is 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' . Q.E.D.

Recall that $\vec{G}(H)$ is the directed graph associated with the hierarchy H (see footnote 9). The next result shows that if H' is a closed subhierarchy of H , then $\vec{G}(H')$ is the subgraph of $\vec{G}(H)$ obtained by dropping edges with nodes that are not in H' .

LEMMA 3: *Fix a hierarchy $H = (N, \geq)$ and a closed subhierarchy $H' = (N', \geq)$. Let \vec{E} be the set of edges in $\vec{G}(H)$. Then $\vec{G}(H') = (N', \vec{E}')$, where $\vec{E}' = \{(n, n') \in \vec{E} \mid n, n' \in N'\}$.*

PROOF OF LEMMA 3: Fix $n, n' \in N'$. We need to show that $(n, n') \in \vec{E}'$ if and only if $(n, n') \in \vec{E}$. If $(n, n') \in \vec{E}$, then n covers n' in H , that is, $n > n'$ and there is no $n'' \in N$ with $n > n'' > n'$. A fortiori, there is no $n'' \in N' \subset N$ with $n > n'' > n'$; hence, n covers n' in H' , so $(n, n') \in \vec{E}'$. If $(n, n') \in \vec{E}'$, then n covers n' in H' . As a result, $n > n'$. If there exists $n'' \in N$ such that $n > n'' > n'$, then, since H' is closed, $n'' \in N'$, so n must not cover n' in H' , a contradiction. As a result, n covers n' in H as well, so $(n, n') \in \vec{E}$. Q.E.D.

The *between set* of (n, n') in $H = (N, \geq)$ is defined as

$$Btw(n, n', H) = \{\hat{n} \in N \mid n \geq \hat{n} \geq n'\}.$$

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

LEMMA 4: *Fix a hierarchy $H = (N, \geq)$ and a subhierarchy $H' = (N', \geq)$. Suppose H' is an MCC of H . Then, either (i) N' is a simple between set in H , or (ii) every cycle in $G(H')$ is a spanning cycle.*

PROOF OF LEMMA 4: For this proof, all between sets are defined relative to H , and we simply write $Btw(n, n')$ for $Btw(n, n', H)$. Similarly, by closed we mean closed in H . Note that by Lemma 3, $\vec{G}(H')$ is the subgraph of $\vec{G}(H)$ obtained by dropping edges with nodes that are not in H' . In particular, H' is cyclic if and only if H contains a cycle whose nodes are in H' . This fact is used freely below.

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 $\vec{G}(H)$. Note that $Btw(n, n')$ is closed, so that these paths are in $\vec{G}((Btw(n, n'), \geq))$ as well, so that $(Btw(n, n'), \geq)$ is cyclic. 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 $\vec{G}(H)$. 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 n to n' in the directed graph $\vec{G}(H)$. Given a path P (either directed or undirected), let N_P denote the set of nodes that appear in P . Since H' is an MCC, $G(H') = (N', E)$ contains a cycle $C = (n_0, e_0, \dots, n_{L-1}, e_{L-1}, n_L)$ where $L > 1$, $n_0 = n_L$. We will argue that (N_C, \geq) 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, \geq) is not closed, in order to reach a contradiction. Given a directed path $(n_0, \vec{e}_0, \dots, n_{L-1}, \vec{e}_{L-1}, n_L)$, its *undirected analogue* is the undirected path $(n_0, e_0, \dots, n_{L-1}, e_{L-1}, n_L)$ where $e_i = \{n_i, n_{i+1}\}$. Say that a directed path P in $\vec{G}(H)$ *only contains edges in C* if every edge in the undirected analogue of P is in C . A directed path P in $\vec{G}(H)$ 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 $i \leq \underline{l} < \bar{l} \leq j$, $(n_{\underline{l}}, n_{\bar{l}})$ is an ECP only if $\underline{l} = i$ and $\bar{l} = j$.

CLAIM 3: Given any two nodes $n, n' \in N'$, if P is the unique directed path in $\vec{G}(H)$ from n to n' , then $N_P = Btw(n, n')$.

PROOF OF CLAIM 3: If there are two non-comparable nodes in $Btw(n, n')$, there would be two distinct directed paths from n to n' . Hence, all nodes in $Btw(n, n')$ are comparable. Therefore, there is a directed path from n to n' whose nodes are $Btw(n, n')$. Since there is a unique directed path from n to n' , the set of nodes in P is $Btw(n, n')$. *Q.E.D.*

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. *Q.E.D.*

Let (n_i, n_j) be a minimally close ECP s.t. $\{n_i, n_{i+1}, \dots, n_j\} \subsetneq N_C$. Let $\bar{n} = \max\{n_i, n_j\}$ and $\underline{n} = \min\{n_i, n_j\}$. Let P^e denote the external directed connection from \bar{n} to \underline{n} . Let \tilde{P} be the undirected path $(n_i, \tilde{e}_i, \dots, \tilde{e}_{j-1}, n_j)$ in $G(H)$ from n_i to n_j in C . Let \tilde{Q} denote the undirected path from n_i to n_j that “goes in the other direction” from \tilde{P} in C , that is, $\tilde{Q} = (n_i, e_{i-1}, n_{i-1}, \dots, e_0, n_0, e_{L-1}, n_{L-1}, \dots, e_j, n_j)$. Let $S = N_{\tilde{P}} \cup Btw(\bar{n}, \underline{n})$.

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 $G((S, \geq))$. One path is \tilde{P} . The other path is the undirected analogue of the external directed connection P^e . Since P^e is external, these two undirected paths must be distinct. *Q.E.D.*

CLAIM 6: S is closed.

PROOF OF CLAIM 6: Let $Y = \bigcup_{n, n' \in N_{\tilde{P}}} Btw(n, n')$. 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, $Btw(\bar{n}, \underline{n}) \subseteq Y$. 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 $(n', n'') = (\bar{n}, \underline{n})$, $n \in Btw(\bar{n}, \underline{n}) \subseteq S$. Suppose instead that $(n', n'') \neq (\bar{n}, \underline{n})$. We will reach a contradiction. Let P denote the directed path from n' to n'' whose nodes include n . Because (\bar{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 \bar{n} and \underline{n} . The sequence of nodes and edges in P between \bar{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 $Y \subseteq S$. *Q.E.D.*

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 $S = N_{\bar{p}} \cup Btw(\bar{n}, \underline{n})$, it must be that $N_{\bar{Q}} \subseteq Btw(\bar{n}, \underline{n})$. All nodes in a between set are comparable, by the hypothesis that directed paths are unique, and so all nodes in $N_{\bar{Q}}$ are comparable. Hence, \tilde{Q} must be the undirected analogue of P^e . This contradicts the hypothesis that (\bar{n}, \underline{n}) is an ECP. Q.E.D.

LEMMA 5: *If the hierarchy $H = (N, \geq)$ is such that every cycle in $G(H)$ is a spanning cycle, then for any pair of nodes $n, n' \in N$, there exist two undirected paths from n to 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 $\{n, n'\}$.*

PROOF OF LEMMA 5: Since there exists a spanning cycle in $G(H)$, for any pair of nodes $n, n' \in N$, there exist two paths $P = (n, e_0^P, n_1^P, \dots, n_{L^P-1}^P, e_{L^P-1}^P, n')$ and $Q = (n, e_0^Q, n_1^Q, \dots, n_{L^Q-1}^Q, e_{L^Q-1}^Q, n')$ in $G(H)$ such that $N = N_P \cup N_Q$ and $E_P \cap E_Q = \emptyset$, where $E_P = \{e_0^P, \dots, e_{L^P-1}^P\}$ and $E_Q = \{e_0^Q, \dots, e_{L^Q-1}^Q\}$. We need to show that $N_P \cap N_Q = \{n, n'\}$. Suppose to the contrary that there exists $\hat{n} \in N_P \cap N_Q$ with $\hat{n} \notin \{n, n'\}$. We know there exist $l^P \in \{1, \dots, L^P - 1\}$ and $l^Q \in \{1, \dots, L^Q - 1\}$ such that $\hat{n} = n_{l^P}^P = n_{l^Q}^Q$. Now, consider the path $(n, e_0^P, n_1^P, \dots, e_{l^P-1}^P, n_{l^P}^P, e_{l^Q-1}^Q, n_{l^Q-1}^Q, e_{l^Q-2}^Q, \dots, n)$ in $G(H)$. Since $E_P \cap E_Q = \emptyset$, this is a well-defined path. But it is a cycle that is not spanning and thus we have reached a contradiction. Q.E.D.

Suppose a hierarchy H has nodes

$$N = \{\bar{n}\} \cup \{\underline{n}\} \bigcup_{l=1}^L \{n_k^l\}_{k=1}^{K_l},$$

with $L \geq 2$, such that for all $l = 1, \dots, L$, (i) $K_l \geq 1$, (ii) $\bar{n} > n_1^l$, (iii) $n_{K_l}^l > \underline{n}$, (iv) n_k^l and n_{k+1}^l are comparable for every $k = 1, \dots, K_l - 1$, 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 6: *If the hierarchy H is a UNP, then H is not universally constructible.*

PROOF OF LEMMA 6: Let $\Omega = \{0, 1\}$. We will present a monotone experiment allocation based on the allocation on the diamond in Section 3 and show that it is not constructible. Consider the experiment allocation given by²¹ $\beta(\underline{n}) = (\frac{1}{2}, (\frac{1}{6}); \frac{1}{2}(\frac{5}{6}))$; $\beta(\bar{n}) = (\frac{3}{8}, (0); \frac{1}{4}, (\frac{1}{2}); \frac{3}{8}, (1))$, $\beta(n_k^1) = (\frac{3}{8}, (0); \frac{1}{8}, (\frac{2}{3}); \frac{1}{2}, (\frac{5}{6}))$ for $k = 1, \dots, K_1$; and $\beta(n_k^l) = (\frac{1}{2}, (\frac{1}{6}); \frac{1}{8}, (\frac{1}{3}); \frac{3}{8}, (1))$ for all $l = 2, \dots, L$ and $k = 1, \dots, K_l$. These are the experiments depicted in Figure 2, where $\beta(\bar{n}) = \tau_A$; $\beta(\underline{n}) = \tau_D$; for all $k = 1, \dots, K_1$, $\beta(n_k^1) = \tau_B$; and for all $l = 2, \dots, L$ and $k = 1, \dots, K_l$, $\beta(n_k^l) = \tau_C$. Monotonicity of β is immediate from the mean-preserving spreads depicted in Figure 2.²²

²¹The experiment is given as a sequence of ordered pairs of a probability and a belief.

²²To see that τ_B is a mean-preserving spread of τ_D , note that we can obtain τ_B from τ_D by spreading the realization $\mu = \frac{5}{6}$ in τ_D to $\{\frac{1}{3}, 1\}$ in τ_B and leaving the realization $\mu = \frac{1}{6}$ in τ_D unchanged. To see that τ_A is a mean-preserving spread of τ_B , note that we can obtain τ_A from τ_B by spreading the realizations $\mu = \frac{1}{6}$ and $\mu = \frac{1}{3}$ in τ_B to $\{0, \frac{1}{2}\}$ in τ_A and leaving the realization $\mu = 1$ in τ_B unchanged. The argument for why $\tau_A \succsim \tau_C \succsim \tau_D$ is symmetric.

Toward a contradiction, suppose σ is monotone signal allocation that induces β . We begin by establishing that it must be the case that, along any undirected path from \bar{n} to \underline{n} , the realized beliefs must be equal across the interior nodes on that path.

CLAIM 7: For all $l = 1, \dots, L$, we have $\tilde{\mu}_{\sigma(n_k^l)} = \tilde{\mu}_{\sigma(n_1^l)}$ for all $k = 1, \dots, K_l$.

PROOF OF CLAIM 7: It suffices to establish that for any $k = 1, \dots, K_l - 1$, we have $\tilde{\mu}_{\sigma(n_{k+1}^l)} = \tilde{\mu}_{\sigma(n_k^l)}$. We know n_k^l and n_{k+1}^l are comparable. Assume that $n_{k+1}^l > n_k^l$; the case where $n_k^l > n_{k+1}^l$ is analogous and omitted. Since σ is monotone, we have that $\sigma(n_{k+1}^l) \supseteq \sigma(n_k^l)$. As the refinement order implies the belief-martingale property, we have that $\mathbb{E}[\tilde{\mu}_{\sigma(n_{k+1}^l)} | \tilde{\mu}_{\sigma(n_k^l)}] = \tilde{\mu}_{\sigma(n_k^l)}$. Now suppose toward contradiction that $\tilde{\mu}_{\sigma(n_{k+1}^l)} \neq \tilde{\mu}_{\sigma(n_k^l)}$. We would conclude that $\tilde{\mu}_{\sigma(n_{k+1}^l)}$ is a strict mean-preserving spread of $\tilde{\mu}_{\sigma(n_k^l)}$.²³ But that would mean that $\tilde{\mu}_{\sigma(n_{k+1}^l)}$ is not equal to $\tilde{\mu}_{\sigma(n_k^l)}$ in distribution and thus that $\langle \sigma(n_{k+1}^l) \rangle \neq \langle \sigma(n_k^l) \rangle$. Q.E.D.

We will now reach a contradiction by establishing that $\Pr(\tilde{\mu}_{\sigma(\bar{n})} = 0 \& \tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6})$ is equal to zero and is strictly bigger than zero.

Step 1: We show that $\Pr(\tilde{\mu}_{\sigma(\bar{n})} = 0 \& \tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6}) = 0$. First, note that $\Pr(\tilde{\mu}_{\sigma(\bar{n})} = 0 | \tilde{\mu}_{\sigma(n_1^1)} = 0) = 1$ since $\mathbb{E}[\tilde{\mu}_{\sigma(\bar{n})} | \tilde{\mu}_{\sigma(n_1^1)} = 0] = 0$ and the support of $\tilde{\mu}_{\sigma(\bar{n})}$ lies above 0. Moreover, since $\Pr(\tilde{\mu}_{\sigma(\bar{n})} = 0) = \frac{3}{8} = \Pr(\tilde{\mu}_{\sigma(n_1^1)} = 0)$, we must also have $\Pr(\tilde{\mu}_{\sigma(n_1^1)} = 0 | \tilde{\mu}_{\sigma(\bar{n})} = 0) = 1$. A similar argument establishes that $\Pr(\tilde{\mu}_{\sigma(n_{K_1}^1)} = \frac{5}{6} | \tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6}) = 1$ and $\Pr(\tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6} | \tilde{\mu}_{\sigma(n_{K_1}^1)} = \frac{5}{6}) = 1$. Claim 7 tells us that $\Pr(\tilde{\mu}_{\sigma(n_1^1)} = 0 \& \tilde{\mu}_{\sigma(n_{K_1}^1)} = \frac{5}{6}) = 0$. Hence, $\Pr(\tilde{\mu}_{\sigma(\bar{n})} = 0 \& \tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6}) = 0$.

Step 2: We show that $\Pr(\tilde{\mu}_{\sigma(\bar{n})} = 0 \& \tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6}) > 0$. It suffices to show that (a) $\Pr(\tilde{\mu}_{\sigma(n_1^1)} = \frac{1}{3} | \tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6}) > 0$, and (b) $\Pr(\tilde{\mu}_{\sigma(\bar{n})} = 0 | \tilde{\mu}_{\sigma(n_1^1)} = \frac{1}{3} \& \tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6}) > 0$.

Arguments analogous to the ones in Step 1 yield $\Pr(\tilde{\mu}_{\sigma(n_{K_2}^2)} = \frac{1}{6} | \tilde{\mu}_{\sigma(\underline{n})} = \frac{1}{6}) = 1$. Because $\text{Supp}(\tilde{\mu}_{\sigma(\underline{n})}) = \{\frac{1}{6}, \frac{5}{6}\}$, this in turn implies that $\Pr(\tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6} | \tilde{\mu}_{\sigma(n_{K_2}^2)} = \frac{1}{6}) = 1$. Thus, Claim 7 tells us that $\Pr(\tilde{\mu}_{\sigma(n_1^1)} = \frac{1}{3} | \tilde{\mu}_{\sigma(n_{K_2}^2)} = \frac{1}{6}) = 1$. Therefore, $\Pr(\tilde{\mu}_{\sigma(n_1^1)} = \frac{1}{3} | \tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6}) = \frac{\Pr(\tilde{\mu}_{\sigma(n_1^1)} = \frac{1}{3})}{\Pr(\tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6})} > 0$, establishing part (a).

Now, note that $\Pr(\tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6} | \tilde{\mu}_{\sigma(n_1^1)} = \frac{1}{3}) = 1$ implies $\Pr(\tilde{\mu}_{\sigma(\bar{n})} = 0 | \tilde{\mu}_{\sigma(n_1^1)} = \frac{1}{3} \& \tilde{\mu}_{\sigma(\underline{n})} = \frac{5}{6}) = \Pr(\tilde{\mu}_{\sigma(\bar{n})} = 0 | \tilde{\mu}_{\sigma(n_1^1)} = \frac{1}{3})$. Moreover, since $\text{Supp}(\tilde{\mu}_{\sigma(\bar{n})}) = \{0, \frac{1}{2}, 1\}$, we have $\mathbb{E}[\tilde{\mu}_{\sigma(\bar{n})} | \tilde{\mu}_{\sigma(n_1^1)} = \frac{1}{3}] = \frac{1}{3} \Rightarrow \Pr(\tilde{\mu}_{\sigma(\bar{n})} = 0 | \tilde{\mu}_{\sigma(n_1^1)} = \frac{1}{3}) > 0$, establishing (b). Q.E.D.

LEMMA 7: The crown is not universally constructible.

PROOF: Consider the experiment allocation given by $\beta(F) = (\frac{1}{8}, (0, \frac{1}{2}); \frac{1}{8}, (\frac{1}{2}, \frac{1}{2}); \frac{1}{8}, (\frac{1}{2}, 0); \frac{1}{8}, (0, 0); \frac{1}{2}, (\frac{1}{4}, \frac{1}{4}))$, $\beta(M) = (\frac{1}{4}, (0, \frac{1}{4}); \frac{1}{4}, (\frac{1}{2}, \frac{1}{4}); \frac{1}{4}, (\frac{1}{4}, 0); \frac{1}{4}, (\frac{1}{4}, \frac{1}{2}))$, $\beta(S) = (\frac{1}{4}, (\frac{1}{4}, 0); \frac{1}{2}, (\frac{1}{4}, \frac{1}{4}); \frac{1}{4}, (\frac{1}{4}, \frac{1}{2}))$, and $\beta(D) = (\frac{1}{4}, (0, \frac{1}{4}); \frac{1}{2}, (\frac{1}{4}, \frac{1}{4}); \frac{1}{4}, (\frac{1}{2}, \frac{1}{4}))$. These experiments are depicted in Figure 5. To see that $\beta(F)$ is a mean-preserving spread of $\beta(S)$,

²³For any two random variables X and Y , if $\mathbb{E}[X|Y] = Y$ and $X \neq Y$, X must be a strict mean-preserving spread of Y .

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

We claim that β is not constructible. As before, consider the joint distribution of beliefs on $\{F, M, S, D\}$ induced by any monotone signal allocation σ , and recall that $\sigma(n) \supseteq \sigma(n')$ implies $\mathbb{E}[\tilde{\mu}_{\sigma(n)} | \tilde{\mu}_{\sigma(n')}] = \tilde{\mu}_{\sigma(n')}$. Specifically, consider the conditional probability of $\tilde{\mu}_{\sigma(S)} = \mu_0$ given $\tilde{\mu}_{\sigma(D)} = \mu_0$. Since $\sigma(M) \supseteq \sigma(S)$, we have that $\tilde{\mu}_{\sigma(S)} = \mu_0 \Leftrightarrow \tilde{\mu}_{\sigma(M)} \in \{(0, \frac{1}{4}), (\frac{1}{2}, \frac{1}{4})\}$. Since $\sigma(M) \supseteq \sigma(D)$, we have that $\tilde{\mu}_{\sigma(D)} = \mu_0 \Leftrightarrow \tilde{\mu}_{\sigma(M)} \in \{(\frac{1}{4}, 0), (\frac{1}{4}, \frac{1}{2})\}$. Hence, the joint probability of $\tilde{\mu}_{\sigma(S)} = \mu_0$ and $\tilde{\mu}_{\sigma(D)} = \mu_0$ must be zero, that is, $\Pr(\tilde{\mu}_{\sigma(S)} = \mu_0 | \tilde{\mu}_{\sigma(D)} = \mu_0) = 0$. But, $\sigma(F) \supseteq \sigma(S)$ implies $\tilde{\mu}_{\sigma(S)} = \mu_0 \Leftrightarrow \tilde{\mu}_{\sigma(F)} = \mu_0$, and $\sigma(F) \supseteq \sigma(D)$ implies $\tilde{\mu}_{\sigma(D)} = \mu_0 \Leftrightarrow \tilde{\mu}_{\sigma(F)} = \mu_0$. Hence, $\Pr(\tilde{\mu}_{\sigma(S)} = \mu_0 | \tilde{\mu}_{\sigma(D)} = \mu_0) = 1$, a contradiction. *Q.E.D.*

LEMMA 8: *If the hierarchy $H = (N, \geq)$ is cyclic and N is a simple between set in H , then H is not universally constructible.*

PROOF OF LEMMA 8: Since $H = (N, \geq)$ is cyclic and N is a simple between set, H must be a UNP. Therefore, Lemma 6 establishes that H is not universally constructible. *Q.E.D.*

LEMMA 9: *If the hierarchy H is cyclic and not a crown, and every cycle in $G(H)$ is a spanning cycle, then H is not universally constructible.*

PROOF OF LEMMA 9: Suppose that $H = (N, \geq)$ is cyclic and not a crown, and that every cycle in $G(H)$ is a spanning cycle. We say an element \bar{n} of N is maximal if there is no $n \in N$ such that $n > \bar{n}$. We say an element \underline{n} of N is minimal if there is no $n \in N$ such that $\underline{n} > n$.

CLAIM 8: *There exist \bar{n} and \underline{n} in N that are maximal and minimal, respectively, such that \bar{n} does not cover \underline{n} .*

PROOF OF CLAIM 8: Suppose toward contradiction that every maximal element covers every minimal element. Hence, $G(H)$ 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 and therefore is cyclic. Moreover, H' is clearly closed, so that by Lemma 3, the nodes in H' are part of a cycle in $G(H)$ as well. Since every cycle in $G(H)$ is a spanning cycle, we must have $N = N'$, and thus H is a crown subhierarchy, and we have reached a contradiction. *Q.E.D.*

By Claim 8, we can find \bar{n} and \underline{n} in N that are maximal and minimal, respectively, such that \bar{n} does not cover \underline{n} . By Lemma 5, there are two distinct undirected paths P and Q in $G(H)$ from \bar{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 $\{\bar{n}, \underline{n}\}$. As a result, H is a UNP. Lemma 6 therefore implies that H is not universally constructible. *Q.E.D.*

PROOF OF THEOREM 1: ONLY IF: Suppose $G(H)$ is not a forest, that is, it contains a cycle. Since N is finite, H contains a subhierarchy $H' = (N', \geq)$ that is an MCC of H . By Lemma 4, either (i) N' is a simple between set in H , or (ii) every cycle in $G(H')$ is a spanning cycle. Consider case (i). Because H' is closed and N' is a simple between set in H , N' is also a simple between set in H' . Thus, Lemma 8 implies that H' is not universally constructible. Now consider case (ii). If H' is not a crown, then Lemma 9 implies that H' is not universally constructible. If H' is a crown, Lemma 7 implies it is not universally constructible. Therefore, by Lemma 2, H is not universally constructible. *Q.E.D.*

REFERENCES

- ARIELI, ITAI, YAKOV BABICHENKO, FEDOR SANDOMIRSKIY, AND OMER TAMUZ (2021): “Feasible Joint Posterior Beliefs,” *Journal of Political Economy*, 129. [2203]
- BERGEMANN, DIRK, AND STEPHEN MORRIS (2016): “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, 11, 487–522. [2202]
- BLACKWELL, DAVID (1953): “Equivalent Comparisons of Experiments,” *The Annals of Mathematical Statistics*, 265–272. [2187,2201]
- BÖRGERS, TILMAN, ANGEL HERNANDO-VECIANA, AND DANIEL KRÄHMER (2013): “When Are Signals Complements or Substitutes?” *Journal of Economic Theory*, 148, 165–195. [2202]
- BRANDENBURGER, ADAM, AND EDDIE DEKEL (1993): “Hierarchies of Beliefs and Common Knowledge,” *Journal of Economic Theory*, 59, 189–198. [2202]
- BROOKS, BENJAMIN, ALEXANDER FRANKEL, AND EMIR KAMENICA (2022): “Comparisons of Signals,” Working paper. [2198]
- CURELLO, GREGORIO, AND LUDVIG SINANDER (2019): “The Lattice Structure of Preference Comparison,” Tech. rep, University of Oxford and Northwestern University. [2201]
- FILL, JAMES ALLEN, AND MOTOYA MACHIDA (2001): “Stochastic Monotonicity and Realizable Monotonicity,” *Annals of Probability*, 938–978. [2202]
- FUCHS, WILLIAM (2007): “Contracting With Repeated Moral Hazard and Private Evaluations,” *American Economic Review*, 97, 1432–1448. [2200]
- GENTZKOW, MATTHEW, AND EMIR KAMENICA (2017): “Bayesian Persuasion With Multiple Senders and Rich Signal Spaces,” *Games and Economic Behavior*, 104, 411–429. [2189,2190,2202]
- GREEN, JERRY R., AND NANCY L. STOKEY (1978): *Two Representations of Information Structures and Their Comparisons*. [2187-2190,2193,2195,2201,2206]
- HARSANYI, JOHN C. (1967): “Games With Incomplete Information Played by “Bayesian” Players, I–III Part I. The Basic Model,” *Management Science*, 14, 159–182. [2202]
- HIRSCH, FRANCIS, CHRISTOPHE PROFETA, BERNARD ROYNETTE, AND MARC YOR (2011): *Peacocks and Associated Martingales, With Explicit Constructions*. Springer Science & Business Media. [2201]
- JEHIEL, PHILIPPE (2015): “Investment Strategy and Selection Bias: An Equilibrium Perspective on Overconfidence.” [2200]
- JUILLET, NICOLAS (2016): “Peacocks Parametrised by a Partially Ordered Set,” in *Séminaire de Probabilités XLVIII*. Springer, 13–32. [2201,2202]
- KAMENICA, EMIR, AND MATTHEW GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615. [2202]
- KERTZ, ROBERT P., AND UWE RÖSLER (2000): “Complete Lattices of Probability Measures With Applications to Martingale Theory,” *Lecture Notes-Monograph Series*, 153–177. [2197]
- LU, LINYUAN (2014): “On Crown-Free Families of Subsets,” *Journal of Combinatorial Theory, Series A*, 126, 216–231. [2201]
- MATHEVET, LAURENT, AND INA TANEVA (2020): “Optimal Information Hierarchies.” [2203]
- MERTENS, JEAN-FRANCOIS, AND SHMUEL ZAMIR (1985): “Formulation of Bayesian Analysis for Games With Incomplete Information,” *International Journal of Game Theory*, 14, 1–29. [2202]
- MOLAVI, POOYA (2021): “Tests of Bayesian Rationality,” arXiv preprint arXiv:2109.07007. [2203]
- MU, XIAOSHENG, LUCIANO POMATTO, PHILIPP STRACK, OMER TAMUZ et al. (2019): “Blackwell Dominance in Large Samples,” arXiv preprint arXiv:1906.02838. [2202]
- MÜLLER, ALFRED, AND MARCO SCARSINI (2006): “Stochastic Order Relations and Lattices of Probability Measures,” *SIAM Journal on Optimization*, 16, 1024–1043. [2197]

- SHMAYA, ERAN, AND LEEAT YARIV (2016): “Experiments on Decisions Under Uncertainty: A Theoretical Framework,” *American Economic Review*, 106, 1775–1801. [2203]
- SMOLIN, ALEX (2017): “Dynamic Evaluation Design,” Tech. rep, University Library of Munich, Germany. [2200]
- STRASSEN, VOLKER (1965): “The Existence of Probability Measures With Given Marginals,” *The Annals of Mathematical Statistics*, 36, 423–439. [2201]
- WILLIAMSON, OLIVER E. (1967): “Hierarchical Control and Optimum Firm Size,” *Journal of Political Economy*, 75, 123–138. [2200]

Co-editor Asher Wolinsky handled this manuscript.

Manuscript received 12 June, 2020; final version accepted 24 January, 2022; available online 8 February, 2022.