

Hypothetical Beliefs Identify Information

Jonathan Libgober *

May 14, 2021

Abstract

After observing the outcome of a Blackwell experiment, a Bayesian decisionmaker can form (a) posterior beliefs over the state, as well as (b) posterior beliefs she would observe any given signal (assuming an independent draw from the same experiment). I call the latter her *contingent hypothetical beliefs*. I show geometrically how contingent hypothetical beliefs relate to information structures. Specifically, the information structure can (generically) be derived by regressing contingent hypothetical beliefs on posterior beliefs over the state. Her prior is the unit eigenvector of a matrix determined from her posterior beliefs over the state and her contingent hypothetical beliefs. Thus, all aspects of a decisionmaker's information acquisition problem can be determined using ex-post data (i.e., beliefs after having received signals). I compare my results to similar ones obtained in cases where information is modeled deterministically; the focus on single-agent stochastic information distinguishes my work.

*University of Southern California. Email: libgober@usc.edu. I thank Juan Carrillo and Xiaosheng Mu for helpful discussions.

1 Introduction

This paper introduces and studies an object I refer to as *contingent hypothetical beliefs*. Economists are by now quite familiar with the Blackwell formulation of information arrival (Blackwell (1953)). In this framework, a decisionmaker starts with some prior beliefs over an uncertain state, observes the outcome of an experiment—which is a (stochastically drawn) signal, say s , whose distribution depends on the state—and updates her beliefs accordingly via Bayes rule. I define her contingent hypothetical beliefs to be her (updated) beliefs that she *could have* seen any particular signal, say \tilde{s} , given that she did in fact observe s . These beliefs represent a probability distribution over signals, and can alternatively be thought of as the distribution over signals the decisionmaker would expect to face given another independent draw from the experiment. As their name suggests, these beliefs are hypothetical—specifically, related to what the decisionmaker could potentially see given an observation determined by the information structure—but contingent on her posterior beliefs after having seen s (as opposed to her prior beliefs).

There are a number of reasons contingent hypothetical beliefs may be of interest. One reason is practical—an analyst may only have (or only be able to reliably obtain) data about a decisionmaker’s uncertainty *after* her beliefs have updated, and thus would be unable to directly determine aspects of the informational environment from a position of ignorance. In such cases, contingent hypothetical beliefs provide additional information—which, as I will show, the analyst would be able to use.¹ Another reason is mathematical, as my results highlight certain additional structure implicit in models of stochastic information arrival. Contingent hypothetical beliefs turns out to be similar to higher-order beliefs (i.e., beliefs about the beliefs of other players) in multi-agent settings. Indeed, higher-order beliefs concern the signals *other* players might have observed; there is a tight connection between updating beliefs to form higher-order beliefs in multi-agent settings, and updating beliefs to form contingent hypothetical beliefs in this paper. The technical structure I highlight is reminiscent of similar (important) restrictions on informational environments in multi-agent settings, assuming deterministic information arrival; I discuss this connection more concretely in Section 5.1.

That said, understanding the structure of Blackwell experiments seems to be valuable in itself, as recent years have seen a surge in interest in using them to describe information arrival and transmission. Since the celebrated work of Kamenica and Gentzkow (2011) who studied them within the context of a communication game, current work has seen their applications grow dramatically. This has included decision theoretic work on identifying information (as in Lu (2019)), as well as multi-agent and mechanism design settings where it is used to motivate the solution concept of Bayes Correlated Equilibrium (see Bergemann and Morris (2016) for more on this connection). This way of modelling information is clearly influential, and its usefulness has likely still not been

¹As an example, imagine an experiment seeking to determine if subjects’ prior beliefs are biased toward extremes. If a subject is told that the probability assigned to some state should initially be 0.7, then she might simply report this information having understood the instructions well—even if she updates beliefs according to a true prior which assigns that state probability 0.9. In this case, my method suggests it would be possible to detect this bias from ex-post data, even though eliciting the actual prior used in this kind of setting would be unreliable.

exhausted despite extensive recent work.

It is also worth emphasizing just how central the ability to perform contingent reasoning is in economic theory, both in strategic settings and in environments with uncertainty. The importance of contingent reasoning in strategic contexts has been recognized since at least [Akerlof \(1970\)](#), who observed that markets can unravel when players condition on what prices would be offered contingent on seller quality. Such analysis assumes individuals rationally consider and define updated beliefs that follow every possible price the seller might charge. By contrast, [Eyster and Rabin \(2005\)](#) propose a solution concept (cursed equilibrium) whereby players update beliefs about *their own* value correctly conditional on their information, but *do not* subsequently update their beliefs about what other players observed.² The contrast between these papers highlights that the importance of contingent reasoning is generally recognized by economists, at least in strategic settings. Though I focus on the natural first case where decisionmakers *are* able to perform contingent reasoning perfectly, it is important to note that experimental work does show that subjects may have *difficulty* in contingent reasoning—for instance, as illustrated by the Winners Curse (as in [Kagel and Levin \(1986\)](#)) or the Monty Hall problem (as in [Miller and Sanjurjo \(2019\)](#)).³ This raises a natural theoretical question: is it sensible to treat individuals who do not correctly form beliefs about what was not (but could have been) seen as otherwise Bayesian? My results suggest the answer is no. Bayesian updating provides dramatic restrictions on which kinds of contingent hypothetical beliefs can emerge. If a decisionmaker does not reason contingently correctly, then there may *no* information structure which rationalizes these contingent hypothetical beliefs together with her beliefs over the state—a point I make explicitly in Section 3.1.

Thus, it seems natural to ask how beliefs about what a decisionmaker could have seen are limited by her beliefs about what she actually did see. This paper answers this question, describing precisely how contingent hypothetical beliefs are limited by the ex-post beliefs induced by an information structure, and how they relate to one another. As to the best of my knowledge, contingent hypothetical beliefs have not been introduced elsewhere, I am not aware of prior work which would directly answer this question.

Specifically, my main results describe when the underlying informational environment can be determined from the beliefs of the decisionmaker following every signal, together with the contingent hypothetical beliefs. The simplest case is when there are at least as many signals as states. In this case, there is a striking way of determining the underlying information structure: Simply “regress” a contingent belief vector (i.e., the probability assigned to possibly having observed a particular signal) on the matrix of beliefs. While specifying these objects appropriately requires some care, this procedure delivers the information structure which generates the signals. The prior also has a geometric interpretation as an eigenvector of a matrix, which comes out of analyzing a martingale condition on beliefs. Provided this regression is possible, and the eigenvector equation from the

²Again, as I discuss in Section 5.1, there is a fundamental connection between higher-order beliefs and contingent hypothetical beliefs. Thus, the contrast between these two papers is very much related to the exercise at hand.

³See also [Martinez-Marquina et al. \(2019\)](#), who show that the presence of uncertainty dramatically increases the difficulty with which subjects are able to reason about contingencies.

martingale condition has a solution for a unit eigenvector, we can back out both the information structure and the prior from the possible beliefs over the states, together with the contingent hypothetical beliefs. I also discuss what it means for these conditions to fail, and what can be done when they do.

If there are more signals than states, then the problem above cannot be solved via the same method. Again, the issue is familiar from linear regression, where this takes the form of an identification problem that emerges if there are more explanatory variables than observations. It turns out that one proposal from statistics for how to address this problem works in our situation as well—I describe a regularization process which essentially allows us to perform the inversion step required by linear regression. The process is known as “ridge regression.” The idea is to add a small perturbation to the singular matrix to avoid the singularity issue that arises with the identification failure, and thus allow the inversion step. If the correct perturbation is considered, then we recover the information structure again (despite the discontinuity in the limit). While this technique is familiar in statistics and econometrics, I suspect it may have more applications in information economics, which I hope future work may be able to exploit.

My analysis therefore shows how intimately related contingent beliefs are to the underlying information structure. While the particular results describe precisely how to recover the information structure and prior from the contingent hypothetical beliefs, the precise geometric relationship between these different objects appears novel. As alluded to above, the idea that the prior can be interpreted as an eigenvector of a matrix determined from the information structure is reminiscent of an important result in Samet (1998b) in the multi-agent context. However, the question of inferring the information structure itself is not considered (although perhaps it is more precise to say, given the assumption of deterministic information, that actually this step is trivialized). In this paper, this step is needed to determine the prior.⁴ In any event, although aspects of the geometric structure of information I highlight may be familiar in other contexts, the extension to Blackwell experiments—and in particular the corresponding regression interpretation of information structures—appears new and is the primary contribution of this paper.

2 Preliminaries

Let Θ denote a finite state space, and let p be a Bayesian decisionmaker’s prior over Θ . Let $\mathcal{I} : \Theta \rightarrow \Delta(S)$ denote a *Blackwell experiment* or *information structure*, where we take $\mathcal{I}(\theta)[s]$ to refer to the probability that the decisionmaker observes s in state θ . I also assume throughout that S is finite. Together with a prior belief over Θ , and given any signal $s \in S$, a decisionmaker can form beliefs over each state $\theta \in \Theta$ via Bayes rule. Let:

⁴As I discuss more thoroughly in Section 5.1, while a stochastic information structure can be modelled as a deterministic information structure, this equivalence requires the ability to expand the state space. This expansion limits the sharpness of my main results, and makes defining the contingent hypothetical beliefs is less straightforward.

$$b_{s,\theta} = \frac{p(\theta)\mathcal{I}(\theta)[s]}{\sum_{\tilde{\theta} \in \Theta} p(\tilde{\theta})\mathcal{I}(\tilde{\theta})[s]}.$$

I denote the $|S|$ -by- $|\Theta|$ matrix of beliefs (over the state) by $B = (b_{s,\theta})_{s \in S, \theta \in \Theta}$. Here, rows are signal observations, and columns are states $\theta \in \Theta$ over which the beliefs are formed. I refer to the matrix B as the *state belief matrix* when necessary to avoid confusion with the contingent hypothetical belief matrix introduced below. I assume that B does not have a zero vector for any column.

In addition, the decisionmaker has the ability to think counterfactually. I let $q_{s,\tilde{s}}$ denote the probability the decisionmaker assigns to the experiment \mathcal{I} generating signal \tilde{s} , if in fact the decisionmaker observed signal s . I let $Q = (q_{s,\tilde{s}})_{s \in S, \tilde{s} \in S}$ denote the corresponding $|S|$ -by- $|S|$ matrix, where the rows are signals the decisionmaker *observes*, and the columns index the decisionmaker's belief that she *could have observed any particular signal* from the set S (so that Q is row-stochastic). I refer to the matrix Q as the *contingent hypothetical belief matrix* (and each probability $q_{s,\tilde{s}}$ as a contingent hypothetical belief).

Note that given a signal s , a decisionmaker with updated beliefs $b_{s,\theta}$ and knowledge of \mathcal{I} can form her contingent hypothetical beliefs. Indeed, while I treat the matrix Q as a primitive, $q_{s,\tilde{s}}$ could be derived from B and \mathcal{I} as follows:

$$q_{s,\tilde{s}} = \sum_{\theta} \mathcal{I}(\theta)[\tilde{s}] b_{s,\theta}$$

I will call an *informational environment* the combination of the information structure \mathcal{I} and the prior p_0 . To help motivate the exercise, one can imagine the case where \mathcal{I} and p_0 define the agent's problem, but B and Q are the relevant objects observed by an analyst. However, I am more generally interested in whether we can recover \mathcal{I} and p_0 from B and Q , as well as how restricted Q is. For the results below, it will be convenient to index the information structure \mathcal{I} so that each row corresponds to the vector $\mathcal{I}(\theta)[\cdot]$ (that is, so that rows are indexed by states).

2.1 Example: Truth or Noise Information

I walk through a simple example to illustrate the key definitions from the previous section. Suppose $\Theta = \{\theta_1, \theta_2, \theta_3\}$, and suppose the decisionmaker uses an initial prior over Θ that assigns probability p_i to state θ_i . Consider the following information structure: With probability $\varepsilon \in (0, 1)$, the decisionmaker observes a “null signal.” With probability $1 - \varepsilon$, the decisionmaker observes the state. Using the above formalism, we can write the state belief matrix as follows:

$$B = \begin{pmatrix} p_1 & p_2 & p_3 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Notice that, consistent with the definitions above, the rows refer to different signals the decision-

maker may observe, and the columns refer to different states.

The contingent hypothetical belief matrix corresponding to this information structure is therefore:

$$Q = \begin{pmatrix} \varepsilon & (1 - \varepsilon)p_1 & (1 - \varepsilon)p_2 & (1 - \varepsilon)p_3 \\ \varepsilon & 1 - \varepsilon & 0 & 0 \\ \varepsilon & 0 & 1 - \varepsilon & 0 \\ \varepsilon & 0 & 0 & 1 - \varepsilon \end{pmatrix}.$$

To understand why the contingent hypothetical belief matrix takes this form, note that if the decisionmaker were to observe no information, then she would still understand this event to be an ε probability occurrence; this corresponds to the entry in the upper left corner. On the other hand, since she obtains no information following the uninformative signal, the probability she would then assign to observing each of the other three signals, *conditional on observing an informative signal*, simply coincides with the prior distribution. Since the probability of observing such a signal in the first place is $1 - \varepsilon$, she therefore assigns $(1 - \varepsilon)p_i$ to the event that she would have seen the signal revealing state θ_i . Following every other signal, while she *would* know the state, she would also understand that there was an ε chance that she would have remained uninformed. Thus, on the one hand, she assigns probability ε to observing the uninformative signal, but also assigns probability 0 to observing any signal that would reveal any other state. Importantly, this matrix is row stochastic, each row itself being a probability distribution over S .

3 Identifying Information if $|S| \geq |\Theta|$

The previous section showed how to construct both kinds of belief matrices from an information structure given a prior belief; the state belief matrix B is computed from Bayes rule, whereas the contingent hypothetical belief matrix additionally can be derived using rules of conditional and total probabilities. A well-known result (originally due to [Aumann and Maschler \(1995\)](#)) states that, given a prior belief and a set of posteriors, there is an essentially unique information structure inducing these beliefs. The main question I study in this paper is how to infer the decisionmaker's information structure and prior using data in the form of the contingent hypothetical belief matrix.

Theorem 1. *For generic state belief matrices B and contingent hypothetical belief matrices Q with $|S| \geq |\Theta|$, \mathcal{I} and p are uniquely identified by B and Q .*

The Theorem actually shows how to construct the information structure from the state belief matrix and the contingent hypothetical belief matrix. Specifically, it shows that the information structure arises from *regressing a given column of the contingent belief matrix on the columns of the state belief matrix*. I illustrate this using the truth-or-noise information structure from the previous section. If $p_1 = p_2 = p_3 = 1/3$, then I compute:

$$(B^T B)^{-1} B^T = \begin{pmatrix} \frac{1}{4} & \frac{11}{12} & -\frac{1}{12} & -\frac{1}{12} \\ \frac{1}{4} & -\frac{1}{12} & \frac{11}{12} & -\frac{1}{12} \\ \frac{1}{4} & -\frac{1}{12} & -\frac{1}{12} & \frac{11}{12} \end{pmatrix}$$

For an arbitrary vector v , the expression $(B^T B)^{-1} B^T v$ is well-known as the ordinary least squares regression of v on B ; that is, it determines the coefficients β in the equation $v = B \cdot \beta$ which provide the “best-fit” (according to the least square error). Also observe that in this example, the rows of $(B^T B)^{-1} B^T$ sum to 1 (which turns out to be a general property). Letting $\mathbf{1}_n$ denote a vector of 1s of length n , since the first column of Q is $\varepsilon \cdot \mathbf{1}_4^T$, we thus have the coefficients corresponding to this regression are $\varepsilon \cdot \mathbf{1}_3^T$. This is exactly the vector of probabilities of observing the null signal in each of the three states. On the other hand, if we were to instead consider the second column of Q , we have:

$$(B^T B)^{-1} B^T \begin{pmatrix} \frac{1-\varepsilon}{3} \\ 1-\varepsilon \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 1-\varepsilon \\ 0 \\ 0 \end{pmatrix},$$

which is precisely the vector of probabilities that the decisionmaker observes the signal saying the state is θ_1 (that is, the probability that this signal is seen in each of the three states).

Notice that Theorem 1, and in particular the regression interpretation, also clarifies *exactly* what is decided by each column of Q . Each column of Q gives a unique column of the matrix determining the information structure \mathcal{I} . As the following section notes, however, not *every* matrix Q will yield an information structure consistent with B .

The condition necessary in order for this procedure to work is that the columns of B are linearly independent, which the proof of Theorem 1 shows generically holds. If it fails, then the matrix $B^T B$ is not invertible.⁵ But even in (non-generic) cases where the columns are linearly dependent, we can still recover the belief matrix by removing the linear dependencies, although the procedure is somewhat more involved. To illustrate, consider the following example with $\Theta = \{\theta_1, \theta_2, \theta_3, \theta_4\}$, where the third column of B is a linear combination of the first two:

$$B = \begin{pmatrix} 2/3 & 0 & 1/3 & 0 \\ 1/3 & 1/3 & 1/3 & 0 \\ 0 & 2/5 & 1/5 & 2/5 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad Q = \begin{pmatrix} 1/2 & 1/2 & 0 & 0 \\ 1/4 & 1/2 & 1/4 & 0 \\ 0 & 3/10 & 1/2 & 1/5 \\ 0 & 0 & 1/2 & 1/2 \end{pmatrix}$$

In this case, $B^T B$ is not invertible, as the linear independence condition is not satisfied; specifically, the third column is $1/2$ times the first column and $1/2$ times the second column. Ideally, we could “remove” the third state responsible for the linear dependencies. Importantly, since the contingent hypothetical belief matrix makes no reference to the underlying states, it would not change if states

⁵Indeed, this condition always *fails* when $|S| > |\Theta|$, though not when $|S| \leq |\Theta|$ as in this section.

were removed, provided the distributions over the signals do not change.

I now show how to remove the state θ_3 , and subsequently interpret the original state space as an auxiliary one where θ_3 is induced with equal probabilities following θ_1 and θ_2 (*after* these are already drawn). After doing this, it will be possible to recover the information structure and prior. We renormalize B so that it does not include θ_3 , supposing that it is instead generated following θ_1 or θ_2 *after* the state is first drawn from $\{\theta_1, \theta_2, \theta_4\}$. This directly gives us an alternative state belief matrix:

$$\tilde{B} = \begin{pmatrix} 1 & 0 & 0 \\ 1/2 & 1/2 & 0 \\ 0 & 3/5 & 2/5 \\ 0 & 0 & 1 \end{pmatrix}.$$

If we were to have started with \tilde{B} , then we could obtain B by considering the case where the state is “flipped” to θ_3 following θ_1 , with probability $1/3$, and the state is “flipped” to θ_3 following θ_2 , again with probability $1/3$ (and never following θ_4). Indeed, the third column of B is the sum of the first two columns, times $1/2$; and the first two columns of B are the same as the first two columns of \tilde{B} , divided by $2/3$ (and $2/3$ is the probability that the state is “unflipped”). This is the sense in which θ_3 is a linear combination of θ_1 and θ_2 . Now, the matrix $\tilde{B}^T \tilde{B}$ is invertible, and regressing each column of Q on \tilde{B} gives an information structure. In this case:

$$(\tilde{B}^T \tilde{B})^{-1} \tilde{B} Q = \begin{pmatrix} 1/2 & 1/2 & 0 & 0 \\ 0 & 1/2 & 1/2 & 0 \\ 0 & 0 & 1/2 & 1/2 \end{pmatrix}.$$

One can check that this information structure generates \tilde{B} and Q , as Theorem 1 suggests it should, using the prior $\mathbb{P}[\theta_1] = \mathbb{P}[\theta_2] = 3/8$ and $\mathbb{P}[\theta_4] = 1/4$.

Now, notice that in the above interpretation, θ_3 is induced with equal probabilities following θ_1 and θ_2 . So consider the following information structure on the original state space $\{\theta_1, \theta_2, \theta_3, \theta_4\}$:

$$\mathcal{I} = \begin{pmatrix} 1/2 & 1/2 & 0 & 0 \\ 0 & 1/2 & 1/2 & 0 \\ 1/4 & 1/2 & 1/4 & 0 \\ 0 & 0 & 1/2 & 1/2 \end{pmatrix}.$$

Where does the signal distribution following state θ_3 come from? The third row of this vector is one half the first row plus one half the second row. In other words, the signal distribution is exactly what it would be if “the state is θ_3 ” is equivalent to “the state is θ_1 with probability $1/2$ and θ_2 with probability $1/2$.” And indeed, one can check that this information structure, under a uniform prior (which, again, is what the prior would be under the specification of how θ_3 is determined from θ_1 and θ_2), generates B .

Having described how to obtain the information structure, we can now ask: What about the

prior? As alluded to above, [Kamenica and Gentzkow \(2011\)](#) and [Aumann and Maschler \(1995\)](#) demonstrated that an information structure is identified by the prior belief and the set of posteriors (which in this case is B). However, this result requires that the prior is in the interior of the convex hull of the posterior beliefs. While the regression procedure finds the information structure, it does not necessarily guarantee that the resulting prior is in the interior of the convex hull of the posterior beliefs.

Addressing this issue yields another insight into the geometry of information structures. Note that the probability of observing any particular signal determined from \mathcal{I} and p :

$$\sum_{\theta} \mathcal{I}(\theta)[s] \cdot p(\theta) = \mathbb{P}[s].$$

We also have the martingale property of beliefs:

$$\sum_s b_{\theta,s} \mathbb{P}[s] = p(\theta).$$

Substituting in for $\mathbb{P}[s]$ and rewriting in matrix form gives the following identity:

$$B^T \mathcal{I}^T p = p \Rightarrow (B^T \mathcal{I}^T - I)p = p.$$

This equation demonstrates that the prior is therefore a *unit eigenvector of the matrix $B^T \mathcal{I}^T$* (or, by taking transposes, a left eigenvector of $\mathcal{I}B$). And in fact, given the previous observation that the information structure \mathcal{I} can be identified from B and Q , this shows that the prior can as well. The proof verifies that indeed this eigenvalue can be guaranteed to exist, via an appeal to the Frobenius-Perron theorem. It is worth emphasizing that the Frobenius-Perron theorem ensures that the eigenvector which satisfies this equation is *unique* (up to scaling), implying that the prior is pinned down as well. The condition needed in order for this to hold is that the matrix of which the prior is a Perron eigenvector is irreducible.⁶ This condition holds for generic B and Q that induce informational environments (generically, all entries of the matrix will be positive).

Lastly, although the above arguments are sufficient to identify the informational environment, we could use an identical argument to determine $\mathbb{P}[s]$, if for whatever reason this were of more interest than $p(\theta)$. Following the exact same calculations as above—except substituting in for $p(\theta)$ instead of $\mathbb{P}[s]$ —we also have that $\mathbb{P}[s]$ is a unit eigenvector of the matrix $\mathcal{I}^T B^T$ (i.e., switching the order in which we multiply the two matrices). As I discuss below in [Section 5.1](#), while the eigenvector interpretation of the prior is similar to one obtained by [Samet \(1998b\)](#), I am not aware of any analogous characterization of the probability of each signal’s observation. This duality between $\mathbb{P}[s]$ and $\mathbb{P}[\theta]$ therefore appears novel.

⁶See the deterministic information structure example in [Section 4.1](#) for a case where this condition fails, and a description of what is identified if so.

3.1 Example: Information from Contingent Beliefs?

A natural question that emerges from the previous discussion is how restricted the set of contingent hypothetical belief matrices are. Can an arbitrary contingent hypothetical belief matrix be consistent with a state belief matrix? The answer to this question turns out to be no. For instance, suppose the belief matrix induced is the following:

$$B = \begin{pmatrix} 1/4 & 3/4 \\ 3/4 & 1/4 \end{pmatrix} \Rightarrow (B^T B)^{-1} B^T = \begin{pmatrix} -\frac{1}{2} & \frac{3}{2} \\ \frac{3}{2} & -\frac{1}{2} \end{pmatrix}.$$

Now, given some candidate $Q = \begin{pmatrix} a & 1-a \\ 1-b & b \end{pmatrix}$, right multiplying by Q gives:

$$\begin{pmatrix} -\frac{a}{2} + \frac{3(1-b)}{2} & -\frac{1-a}{2} + \frac{3b}{2} \\ \frac{3a}{2} - \frac{1-b}{2} & \frac{3(1-a)}{2} - \frac{b}{2} \end{pmatrix}.$$

For an arbitrary a, b , it may be that this yields an information structure, yet does not necessarily yield one that can induce B given any prior p . An immediate corollary of [Kamenica and Gentzkow \(2011\)](#) shows that the prior together with B pins down the information structure; from this, the matrix Q can always be inferred. It follows that:

Proposition 1. *For every belief matrix B with linearly independent columns, the dimensionality of Q yielding B is $|\Theta| - 1$.*

Given this result, in the previous example, we see that given a possibly valid a , there is a unique value of b which corresponds to a valid information structure. [Figure 1](#) shows how, given the belief matrix B , which a and b choices correspond to a fixed feasible prior p . For instance, we see that $a = b = 5/8$ is the solution when $p = 1/2$. The choice of $a = b = 9/16$ is therefore invalid; nevertheless, for these choices, we have:

$$(B^T B)^{-1} B^T Q = \begin{pmatrix} 3/8 & 5/8 \\ 5/8 & 3/8 \end{pmatrix}$$

Upon inspection, we see that this *is* indeed a perfectly valid information structure \mathcal{I} , and in fact one that is symmetric. But, it is also straightforward to see that it cannot induce the belief matrix B for any prior; indeed, since B is symmetric as well, symmetry would require $p_0 = 1/2$, which in turn would suggest distinct beliefs given the information structure than those given by B .

If we take any square matrix Q whose rows sum to 1, then $(B^T B)^{-1} B^T Q$ will be a matrix whose rows sum to 1, and thus be a valid information *structure*.⁷ The dimension of the set of row-stochastic $|S|$ -by- $|S|$ matrices is $|S|(|S| - 1)$. On the other hand, [Proposition 1](#) suggests a smaller dimensionality on the set of matrices Q which yield valid informational *environments* inducing B . For instance, in the previous example we can choose a to be anything within some interval, but this will pin down b exactly, so that the dimension of the set of Q inducing B is one; by contrast,

⁷See the Appendix for a proof of this Claim.

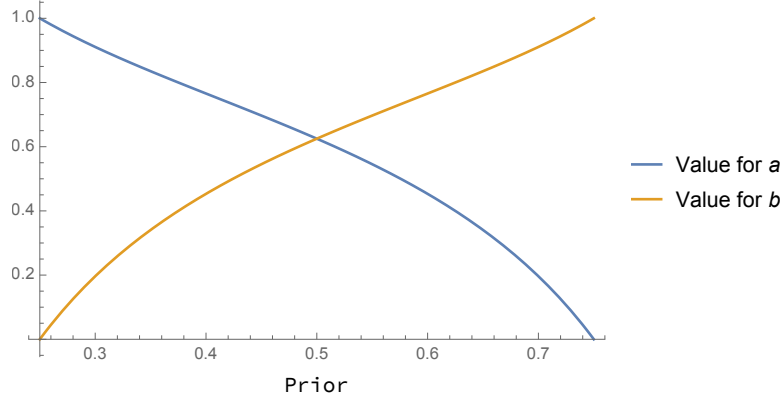


Figure 1: Value for a and b in the example which yield a valid contingent hypothetical belief matrix, given a prior probability of state 1 equal to $p \in [1/4, 3/4]$.

the set of 2-by-2 row stochastic matrices is two dimensional. Therefore, while this procedure does determine an information structure for generic values of Q , it does not necessarily generate an information environment inducing these matrices, because it may be that no prior exists which does so.

4 The $|\Theta| > |S|$ case

The regression characterizations no longer apply in the case where there are more states than signals, in the same way that an Ordinary Least Squares regression requires more observations than covariates in order to yield an identified solution (i.e., B must have more rows than columns, in addition to having no linear dependencies). Appealing to the regression interpretation, however, allows us to appeal to a suggestion from statistics on how to circumvent the problem.

To solve for the information structure, as before, we first pre-multiply the matrix equation for the contingent hypothetical beliefs by the state belief matrix B :

$$B^T Q = B^T B \mathcal{I}$$

If $|S| < |\Theta|$, $B^T B$ is not invertible. However, for all but at most finitely many $\lambda > 0$, the matrix $B^T B + \lambda I$ is invertible, since the determinant is an n -degree polynomial in λ (and since any matrix with non-zero determinant is invertible). This implies:

$$B^T Q + \lambda \mathcal{I} = (B^T B + \lambda I) \mathcal{I} \Rightarrow (B^T B + \lambda I)^{-1} B^T Q = \mathcal{I} (1 - (B^T B + \lambda I)^{-1} \lambda I) \quad (1)$$

Since this holds for all λ , we therefore will be able to recover \mathcal{I} by considering the limit $\lambda \rightarrow 0$; a necessary and sufficient condition for this to converge to \mathcal{I} is:

$$(B^T B + \lambda I)^{-1} \lambda I \rightarrow 0. \quad (2)$$

In statistics, the idea of adding λI in order to be able to invert $B^T B$ is used to motivate *ridge regression*; in cases where more general perturbations are considered, this process is also known as *Tikhonov regularization*. In statistical applications where this method is used, taking λ too small is often undesirable.⁸ Here, however, the $\lambda \rightarrow 0$ limit will prove to be useful.

Unfortunately, the limit defined in (2) need not always hold, although we can still use these ideas to recover the information structure even in that case as well. Before doing this, however, I walk through examples illustrating the key ideas.

4.1 Examples

Consider the following state belief matrix and contingent hypothetical belief matrix, with $|\Theta| = 3$ and $|S| = 2$:

$$B = \begin{pmatrix} 2/3 & 1/3 & 0 \\ 0 & 1/3 & 2/3 \end{pmatrix}, Q = \begin{pmatrix} 5/6 & 1/6 \\ 1/6 & 5/6 \end{pmatrix}.$$

In this example, it turns out Tikhonov regularization as described works to infer the information structure. We emphasize that $\lim_{\lambda \rightarrow 0} (B^T B + \lambda I)^{-1}$ is not well-defined since B is not full rank. However, for all $\lambda > 0$, we have:

$$(B^T B + \lambda I)^{-1} = \begin{pmatrix} \frac{1}{6\lambda} + \frac{1}{2+3\lambda} + \frac{9}{8+18\lambda} & -\frac{2}{6\lambda+9\lambda^2} & \frac{4}{24\lambda+90\lambda^2+81\lambda^3} \\ -\frac{2}{6\lambda+9\lambda^2} & \frac{2}{3\lambda} + \frac{1}{2+3\lambda} & -\frac{2}{6\lambda+9\lambda^2} \\ \frac{4}{24\lambda+90\lambda^2+81\lambda^3} & -\frac{2}{6\lambda+9\lambda^2} & \frac{1}{6\lambda} + \frac{1}{2+3\lambda} + \frac{9}{8+18\lambda} \end{pmatrix}$$

In this case, we have $B^T B$ is non-invertible, but $B^T B + \lambda I$ is, for all $\lambda > 0$. On the other hand, every term in this matrix approaches ∞ as $\lambda \rightarrow 0$. But post-multiplying by B^T gives

$$(B^T B + \lambda I)^{-1} B^T = \begin{pmatrix} \frac{1}{2+3\lambda} + \frac{3}{4+9\lambda} & -\frac{2}{8+30\lambda+27\lambda^2} \\ \frac{1}{2+3\lambda} & \frac{1}{2+3\lambda} \\ -\frac{2}{8+30\lambda+27\lambda^2} & \frac{1}{2+3\lambda} + \frac{3}{4+9\lambda} \end{pmatrix}$$

The entries in this expression do not diverge and this converges to a finite matrix. We can therefore compute:

$$\lim_{\lambda \rightarrow 0} (B^T B + \lambda I)^{-1} B^T Q = \begin{pmatrix} 1 & 0 \\ 1/2 & 1/2 \\ 0 & 1 \end{pmatrix}.$$

It is straightforward to verify that indeed this information structure induces this state belief matrix and contingent hypothetical belief matrix when s_1 is observed with probability 1 in state θ_1 and 1/2 in state θ_2 , and signal s_2 is observed with probability 1/2 in state θ_2 and 1 in state θ_3 .

Our second example both illustrates a case where the regularization method does *not* work, as

⁸See [van Wieringen \(2015\)](#) for more on what guides the choice of λ in statistical applications.

well as serving as a prelude to some discussion in Section 5.1. Specifically, we consider the case of a *deterministic* information structure: In these models, signals are associated with elements of a partition of the state space. Suppose $\Theta = \{\theta_1, \theta_2, \theta_3, \theta_4\}$, and consider an information structure where the decisionmaker observes which element of $\mathcal{P} = \{\{\theta_1\}, \{\theta_2, \theta_3\}, \{\theta_4\}\}$ the state belongs to. If p_i is the prior probability over state θ_i , then this corresponds to the following state belief matrix and hypothetical belief matrix:

$$B = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{p_2}{p_2+p_3} & \frac{p_3}{p_2+p_3} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad Q = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

Of course, in this example our main exercise is fairly straightforward, but understanding what the procedure produces in this case is instructive. For this example, we compute:

$$\lim_{\lambda \rightarrow 0} (B^T B + \lambda I)^{-1} \lambda I = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \frac{p_3(p_3-p_2)}{p_2^2+p_3^2} & 0 \\ 0 & \frac{p_2(p_2-p_3)}{p_2^2+p_3^2} & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

This matrix converges to 0 only when $p_2 = p_3$. And indeed, we see that the procedure fails to produce an information structure:

$$\lim_{\lambda \rightarrow 0} (B^T B + \lambda I)^{-1} B^T Q = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{p_2(p_2+p_3)}{p_2^2+p_3^2} & 0 \\ 0 & \frac{p_3(p_2+p_3)}{p_2^2+p_3^2} & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

What went wrong? Notice that in the above derivation, adding λI to $B^T B$ was only one way to ensure we the inversion step would be possible. And indeed, the equation $Q = BI$ *should* have multiple solutions in the case of $|\Theta| > |S|$, and so the limit only considers one of them.

Note that in this case, a different regularization would deliver the information structure. For instance, one can compute that:

$$\lim_{\lambda \rightarrow 0} \left(B^T B + \lambda \overbrace{\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{p_3}{p_2} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}}{:=M} \right)^{-1} B^T Q = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

which is indeed the deterministic information structure in this example. One can derive this expression by following the same steps as outlined in Equation (1), but considering a different perturbation;

namely, using λM instead of λI .

4.2 Identifying Information

Our main result from this section shows that, even if (2) does not hold, the result of this limit is still useful as it allows us to identify a linear subspace in which the information structure must belong:

Theorem 2. *Suppose $|\Theta| > |S|$. The matrix:*

$$\tilde{\mathcal{I}} = \lim_{\lambda \rightarrow 0} (B^T B + \lambda I)^{-1} B^T Q,$$

exists and is well-defined. Furthermore, let v_1, v_2, \dots, v_k be a basis for the null-space of B . Then $\mathcal{I}(\cdot)[s_i] = \tilde{\mathcal{I}}(\cdot)[s_i] + \sum_j \alpha_j v_j$.

This theorem suggests a way to determine the information structure when there are more states than signals. Of course, when \mathcal{I} has more signals than states and B is full rank, then $B^T B$ is invertible, and the theorem remains true as well—although in this case, considering the limits is unnecessary.

As discussed above, a necessary and sufficient condition for $\mathcal{I} = \tilde{\mathcal{I}}$ is that (2) holds. In fact, more generally, the same argument as outlined above shows, for an arbitrary Θ -by- Θ matrix M , if $\mathcal{I} = \lim_{\lambda \rightarrow 0} (B^T B + \lambda M)^{-1} B^T Q$ exists, then:

$$(B^T B + \lambda M)^{-1} \lambda M \mathcal{I} \rightarrow 0$$

implies that \mathcal{I} induces B and Q . Or, put differently, if one finds that:

$$(B^T B + \lambda M)^{-1} \lambda M (B^T B + \lambda M)^{-1} B^T Q \rightarrow 0,$$

then the analyst can be assured that \mathcal{I} is valid and generates B and Q .

On the other hand, the following example shows that Theorem 2 can still be used to determine the information structure in the case where $|\Theta| > |S|$. Suppose:

$$B = \begin{pmatrix} \frac{2}{3} & \frac{1}{3} & 0 \\ \frac{4}{9} & \frac{1}{9} & \frac{4}{9} \end{pmatrix}, \quad Q = \begin{pmatrix} \frac{7}{18} & \frac{11}{18} \\ \frac{11}{54} & \frac{43}{54} \end{pmatrix}.$$

In this case, we compute $\tilde{\mathcal{I}}$ to be:

$$\begin{pmatrix} \frac{29}{63} & \frac{52}{63} \\ \frac{31}{126} & \frac{23}{126} \\ -\frac{4}{63} & \frac{58}{63} \end{pmatrix}$$

Obviously, $\tilde{\mathcal{I}}$ is not an information structure, as each row violates the two requirements to be probability distributions: non-negative entries and summing to 1. The nullspace of B is spanned

by a single vector, $(-2, 4, 1)$. We add a multiple of this vector to the first column of $\tilde{\mathcal{I}}$ to make sure all entries are non-negative. This yields:

$$\begin{pmatrix} \frac{1}{3} & \frac{52}{63} \\ \frac{1}{2} & \frac{23}{126} \\ 0 & \frac{58}{63} \end{pmatrix}$$

Adding another multiple of $(-2, 4, 1)$ to the second column allows each row to sum to 1, and yields:

$$\begin{pmatrix} \frac{1}{3} & \frac{2}{3} \\ \frac{1}{2} & \frac{1}{2} \\ 0 & 1 \end{pmatrix}$$

which indeed is an information structure inducing these belief matrices.

I briefly mention that the interpretation of the prior as an eigenvector of $B^T \mathcal{I}^T$ remains in this setting as well, as this argument does not rely at all upon the cardinality of either the signal space nor the message space. However, I emphasize that the Frobenius-Perron theorem does require that this matrix is positive; notably, this is *not* satisfied for the deterministic information structure case. And indeed, we do not have a unique prior identified by B and Q in this case—while we can determine the prior beliefs “within a signal,” in the deterministic information structure case, it is not possible to determine the relative probability across different partition elements.⁹

5 Discussion

5.1 Deterministic vs. Stochastic Information

The focus on contingent reasoning embedded within information structures is itself not new, but to the best of my knowledge it has so far been focused on so-called “deterministic” models of information arrival with multiple agents. An important property of Blackwell experiments is that information is stochastically generated as a function of the state. An alternative is to treat signals as *deterministic* functions of the state, and to view an information structure as partitions of a state space (with signals being elements of the partition). It is a common exercise to study the *hierarchy of beliefs* using this formulation; that is, beliefs over other players beliefs, and belief over these beliefs over beliefs, and so on.

As discussed above, deterministic information arrival is immediately a special case of stochastic information arrival. In fact, an information structure is deterministic if and only if Q is the identity (a claim I show in the Appendix). Note that in this case, we will have $|S| > |\Theta|$ whenever the information structure does not reveal the state.

However, it also turns out that stochastic information arrival *can also* be modelled as a special

⁹More generally, as illustrated by this example, if $B^T \mathcal{I}^T$ is not irreducible, then while we can use the Frobenius-Perron theorem to determine a unique prior within each irreducible class, we cannot determine the *relative* prior probabilities across classes.

case of deterministic information arrival, although to do so requires replacing the state space. Green and Stokey (1978) articulated this method clearly; if one views the state itself as belonging to the set $\Theta \times [0, 1]$, instead of Θ , then any stochastic information structure can be represented as a deterministic one, assuming a uniform prior on this second dimension. While this transformation does not satisfy the finiteness requirement I impose on the state space, this limitation is minor.¹⁰ A more substantive restriction, however, is that under this relabeling, one *always* has more signals than states, by construction. In principle, one could define contingent hypothetical beliefs in this framework as well; in this case, these are decisionmaker’s beliefs that she would see a given signal, assuming the same state but an independent draw from the second dimension. If one imagined this second draw were observed by a different agent, then the beliefs about the signals of the others would be a *higher-order belief*. On the other hand, to define the state belief matrix in this case, one would have to “integrate out” the second dimension, something which would not be meaningful if we did not start out with a state space that already had the implied structure of a Blackwell experiment.

Focusing on questions related to the existence of a common prior—and not the identification of the information structure—a similar construction to the one I proposed in this paper was uncovered by Samet (1998b), and further explored by Samet (1998a) and Golub and Morris (2017). Samet (1998b) characterized a *common* prior as the eigenvector of a stochastic transition matrix obtained from an information structure involving multiple agents (specifically, Propositions 3 and 5).¹¹ These results provide an *interim* characterization of the common prior assumption, i.e., after signals have been observed. This paper’s motivation is similar, though I am not only interested in identifying the prior from interim beliefs but also the information structure. On the other hand, once an information structure is assumed to be deterministic, it is immediate to then find the information structure—simply group any signals where beliefs are constant.

In addition, note that, as discussed above, the application of the Frobenius-Perron theorem, used to determine the prior beliefs, would not work in any single agent case with a deterministic information structure. That is, the prior is generally not pinned down by only a single agent’s (non-degenerate) deterministic information structure. In the multiple agent case, however, this limitation can be overcome; again, Samet (1998b)’s result is that a *common* prior is uniquely pinned down as the eigenvector of a *matrix derived by multiplying the belief matrices of each agent regarding the signals received by the other agents*. Golub and Morris (2017) explore this idea further, roughly speaking using similar ideas to move from studying “beliefs about beliefs” to “expectations about expectations”—and subsequently characterizing implications, for instance, of assuming these expectations do not depend on the order agents are considered. Their analysis also involves an explicit application of the Frobenius-Perron theorem, in order to characterize the “consensus expectation”

¹⁰For instance, it would not emerge if one could somehow use $\{0, 1/n, \dots, (n-1)/n, 1\}$ instead of $[0, 1]$, even though this could only generate certain stochastic information structures.

¹¹Specifically, the existence of the common prior is equivalent to this eigenvector being independent of the order in which agents are considered in defining this stochastic transition matrix. See Hellman (2011) for a generalization of this result to infinite spaces.

in a model with interacting agents. Notice that in any deterministic information structure case, many different priors can typically generate the same state belief matrix B (for a single agent). The existence of a common prior is much more restrictive, as the same prior would need to induce possibly different belief matrices B .¹²

By contrast, in a model of stochastic information arrival, the prior may be pinned down even with only a single agent. In addition, the question of how to infer the information structure is trivialized in the case of deterministic information structures (although as I discussed, one could still use my method even there), and so it is not clear that the regression interpretation achieves much in these cases, unlike in the stochastic information case. Still, the question of inferring the information of agents appears practically relevant, and as discussed above, in many cases this may be more naturally modelled using the Blackwell (stochastic) formulation.

5.2 Other Relevant Literature

It is additionally worth pointing out connections to a few other relevant literatures this paper is related to. The first is on the general agenda of determining when it is possible (and if so, how) to identify the parameters of an agent’s decision problem as a function of observables. The problem of identification is one that is commonly studied in econometrics, as well as in decision theory. For the question studied here, the closest paper is [Lu \(2019\)](#), who asks when it is possible to identify a decisionmaker’s utility function and information from choice data, if their stochastic choice is generated by a distribution over posterior beliefs; he shows that it is possible provided the analyst has data from two treatments where the agent is better informed than one or the other, and given data from only one treatment if the signal structure is known. The question in this paper is in a similar spirit, although the objects of interest differ. Of course, the relevance of the results in either paper depend on practical considerations, and hence a productive area for future work may be to determine when identification is possible given limitations of real-world data.

The second idea explored here relates to the use of linear projection methods. It turned out that these techniques were natural in the setting of this paper, since hypothetical beliefs themselves involve a linear relationship between the information structure and beliefs. More generally, linear structure (in various forms) often naturally lends itself to a substantial degree of tractability in a number of settings of economic relevance. For instance, certain problems related to the elicitation of information via scoring rules also inherit this structure, and linear projection methods have proved useful there as well (see, for instance, [Lambert \(2019\)](#) and [Ball \(2021\)](#)). Here, the explicit connection to linear regression was useful in suggesting how to tackle the problem of $|\Theta| > |S|$. However, the use of Tikhinov regularization in this paper appears to be relatively unexplored in microeconomic theory. Since the technical ideas developed here have been used in other settings, one can imagine using further tools developed in statistics to extend the boundary of problems

¹²Several other papers have used either the Markov chain interpretation of interim beliefs introduced by [Samet \(1998b\)](#), or properties implied by the stationary distribution characterization of the common prior; see, for instance, [Morris and Shin \(2002\)](#), [Cripps et al. \(2008\)](#), or [Angeletos and Lian \(2018\)](#).

where linearity provides a significant source of tractability.

5.3 Final Comments

The review of the literature suggests one reason the results of this paper may be of interest, namely in terms of showing linear algebraic characterizations of informational environments in deterministic, multi-agent settings also apply to stochastic, single-agent settings. The objects differ slightly and use the introduction of contingent hypothetical beliefs, which function similarly to higher-order beliefs (since they represent “beliefs about beliefs”). On the other hand, they are quite a bit more restricted than higher-order beliefs, since the universal type space as a mathematical object involves significant added richness (e.g., it may be that two information structures agree in terms of first and second order beliefs, but not on beliefs of order higher than that). Studying contingent hypothetical beliefs in a multi-agent context is a distinct problem, but one that is likely worth studying further in future work.

On the other hand, the primary motivation for the introduction of contingent beliefs is in their importance in describing decision-making in the presence of learning. In this paper, I have sought to articulate how the structure of these beliefs is restricted in a classic Bayesian model. One conclusion a reader may reach from this is that they are too restricted, and that perhaps it is excessively demanding to assume that decisionmakers have such a high degree of consistency between the set of possible beliefs over the state and the set of contingent hypothetical beliefs. As this paper shows, however, it is impossible to have this view without also being critical of the Blackwell formulation of information arrival (i.e., where decisionmakers update beliefs as Bayesian following an observation from a known Blackwell experiment), as these are ultimately equivalent.

References

- AKERLOF, G. (1970): “The Market for “Lemons”: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, 84, 488–500.
- ANGELETOS, G.-M. AND C. LIAN (2018): “Forward Guidance without Common Knowledge,” *American Economic Review*, 108, 2477–2512.
- AUMANN, R. J. AND M. MASCHLER (1995): *Repeated Games with Incomplete Information*, MIT Press.
- BALL, I. (2021): “Scoring Strategic Agents,” *Working Paper*.
- BERGEMANN, D. AND S. MORRIS (2016): “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, 11, 487–522.
- BLACKWELL, D. (1953): “Equivalent Comparison of Experiments,” *Annals of Mathematical Statistics*, 24, 265–272.

- CRIPPS, M., J. ELY, G. MAILATH, AND L. SAMUELSON (2008): “Common Learning,” *Econometrica*, 76, 909–933.
- EYSTER, E. AND M. RABIN (2005): “Cursed Equilibrium,” *Econometrica*, 73, 1623–1672.
- GOLUB, B. AND S. MORRIS (2017): “Higher-Order Expectations,” *Working Paper, Northwestern University and MIT*.
- GREEN, J. AND N. STOKEY (1978): “Two Representations of Information Structures and Their Comparisons,” *Unpublished*.
- HELLMAN, Z. (2011): “Iterated Expectations, Compact Spaces and Common Priors,” *Games and Economic Behavior*, 72, 163–171.
- KAGEL, J. H. AND D. LEVIN (1986): “The Winner’s Curse and Public Information in Common Value Auctions,” *American Economic Review*, 76, 894–920.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- LAMBERT, N. (2019): “Elicitation and Evaluation of Statistical Forecasts,” *Working Paper*.
- LU, J. (2019): “Bayesian Identification: A Theory for State-Dependent Utilities,” *American Economic Review*, 101, 2590–2615.
- MARTINEZ-MARQUINA, A., M. NIEDERLE, AND E. VESPA (2019): “Failures in Contingent Reasoning: The Role of Uncertainty,” *American Economic Review*, 109, 3437–3474.
- MEYER, C. D. (1995): *Matrix Analysis and Applied Linear Algebra*, SIAM.
- MILLER, J. AND A. SANJURJO (2019): “A Bridge from Monty Hall to the Hot Hand: The Principle of Restricted Choice,” *Journal of Economic Perspectives*, 33, 144–162.
- MORRIS, S. AND H. S. SHIN (2002): “Social Value of Public Information,” *American Economic Review*, 92, 1521–1534.
- SAMET, D. (1998a): “Common Priors and Separation of Convex Sets,” *Games and Economic Behavior*, 24, 172–174.
- (1998b): “Iterated Expectations and Common Priors,” *Games and Economic Behavior*, 24, 131–141.
- VAN WIERINGEN, W. N. (2015): *Lecture Notes on Ridge Regression*, ArXiv Preprint.

Proof of Theorem 1. We first note that B is full column rank, for generic belief matrices B with $|\Theta| \leq |S|$ (in other words, when there are more rows than columns). Indeed, the set of all possible belief matrices is of dimension $|S| \times (|\Theta| - 1)$ (with one degree of freedom lost for every row, since every row is restricted to sum to 1); in general, the space spanned by k belief vectors, restricted to sum to 1, is $(k - 1) \times |S|$. Hence a generic belief matrices is full rank, since the set of belief matrices which fail to satisfy this belong to a lower dimension subspace. Given this observation, we have that B is of rank equal to $|\Theta|$. On the other hand, the rank of $B^T B$ is equal to the rank of B , and therefore $B^T B$ has full rank. Since a matrix is invertible if and only if it has full rank, we have that $B^T B$ is invertible.

The rest of the first part of the theorem follows the arguments as laid in the main text. Writing the information structure so that rows are states and columns are signals, the definition of hypothetical beliefs matrix tells us that:

$$Q = B \cdot \mathcal{I},$$

since the i th row and j th column of Q is the inner product of the i th row of B (since rows of B index signals) and the j th column of \mathcal{I} (using the convention that columns index signals). Given this expression, and using that $(B^T B)^{-1}$ is invertible, the solution for \mathcal{I} comes from left-multiplying both sides by B^T and then left multiplying by $(B^T B)^{-1}$.

Next, we show that the prior is identified. As shown in the main text, the prior is a unit eigenvector of the matrix $B^T \mathcal{I}^T$, and therefore a unit eigenvector of $B^T Q^T B (B^T B)^{-1}$. We show that this matrix is *always* row-stochastic. Recall that $\mathbf{1}^T B^T$ and $\mathbf{1}^T Q^T$ are both $\mathbf{1}$, since both B and Q are row-stochastic (so that the transposes are column-stochastic). Therefore, we have that:

$$\mathbf{1}^T B^T Q^T B (B^T B)^{-1} = \mathbf{1}^T Q^T B (B^T B)^{-1} = \mathbf{1}^T B (B^T B)^{-1}$$

Taking the transpose of this expression gives:

$$(B^T B)^{-1} B^T \mathbf{1}.$$

Now, recall the regression interpretation of the linear mapping $(B^T B)^{-1} B^T$; when applied to a vector, it gives the coefficient of the regression of the unit vector onto the columns of B . However, the columns of B sum to 1. Therefore, to write the vector $\mathbf{1}$ as a linear combination of the columns of B , we need only write it using coefficients equal to 1, and hence this expression is itself a unit vector.

Hence, the Frobenius-Perron theorem holds provided the matrix $B^T \mathcal{I}^T$ is irreducible. This property holds generically; indeed, the entries of $B^T \mathcal{I}^T$ are generically positive, and all such matrices are irreducible. This theorem therefore yields the existence of a Perron eigenvector, which is positive and sums to 1. While the argument in the main text shows that being a Perron eigenvector is a *necessary* condition for the prior, the Frobenius-Perron theorem implies that this vector is unique, and therefore this condition is also sufficient. As a result, the prior is additionally identified from

B and Q , in addition to the information structure, as desired. \square

Proof of Proposition 1. By Proposition 1 of [Kamenica and Gentzkow \(2011\)](#), given any belief matrix B and prior p_0 , there exists an information structure \mathcal{I} inducing this belief matrix.¹³ On the other hand, Theorem 1 shows that *any* vector v of length $|S|$ yields a vector of length Θ when considering $(B^T B)^{-1} B^T v$. Thus, the set of information structures is spanned by the set of Q that emerge in some informational environment. Putting the previous observations together, the set of Q which induce an informational environment given a belief matrix B is isomorphic to the set of priors inducing B , which has dimensionality equal to $|\Theta| - 1$, for any belief matrix satisfying the linear independence condition. \square

Proof of the claim that rows of $(B^T B)^{-1} B^T Q$ sum to 1 if Q has rows that sum to 1. The sum of the rows of this matrix is given by right multiplying by $\mathbf{1}_{|S|}$, a vector of length $|S|$ which is all 1s. If Q has rows which sum to 1, then $(B^T B)^{-1} B^T Q \cdot \mathbf{1}_{|S|} = (B^T B)^{-1} B^T \mathbf{1}_{|S|}$. On the other hand, recall that this expression is also the coefficients β_1, \dots, β_n solving:

$$\mathbf{1}_{|S|} = \sum_{i=1}^n \beta_i b_i,$$

where b_i is the i th column of the belief matrix B_i . While there is a unique set of coefficients solving this equation, we also have that the columns of a belief matrix sum to 1. Hence $\beta_1 = \dots = \beta_n = 1$ is the solution. We thus conclude that $(B^T B)^{-1} B^T Q \cdot \mathbf{1}_{|S|}$ is a vector of 1s, as claimed. \square

Proof of Theorem 2. That there is convergence as $\lambda \rightarrow 0$ appears known, although I provide a proof for completeness as other proofs I have found require significant detours.¹⁴ We first determine the rate at which the determinant of $B^T B + \lambda I$ tends to 0 as $\lambda \rightarrow 0$. Note that, by the Matrix Determinant Lemma (see (6.2.3) of [Meyer \(1995\)](#) for a version of this result), we have:

$$\det\left(\frac{1}{\lambda} B^T B + I\right) = \det\left(I_{|S|} + \frac{1}{\lambda} B B^T\right).$$

Note that the matrix involved in the left-hand side of this equation is $|\Theta| - by - |\Theta|$ and the matrix involved in the right hand side of this equation is $|S| - by - |S|$. We therefore have, multiplying through by $\lambda^{|\Theta|}$ and using that $\det(cA) = c^n \det(A)$ for $c \in \mathbb{R}$ and A an n -by- n matrix,

$$\det(B^T B + \lambda I) = \det(\lambda^{|\Theta|/|S|} I_{|S|} + \lambda^{(|\Theta|-|S|)/|S|} B B^T).$$

Note that this determinant is a polynomial in λ which evaluates to 0 at $\lambda = 0$, and hence this approaches 0 at a rate equal to the rate of the smallest term in this polynomial. We claim the

¹³Their proof is constructive; in our notation, one can set $\mathcal{I}(\theta)[s] = B_{s,\theta} \mathbb{P}[s]/p_0[\theta]$. Now, note that $\mathbb{P}[s]$ is a left unit eigenvector of the matrix Q . Note, however, that $\mathbb{P}[s]$ would be derived from Q in this paper, and not the prior p_0 and the posterior beliefs, as in theirs.

¹⁴As per [van Wieringen \(2015\)](#), the limit of the ridge estimator as $\lambda \rightarrow 0$ is precisely the least square estimate of smallest norm; showing this, however, requires a significant detour in defining ridge estimators. Note that [van Wieringen \(2015\)](#) also shows that multiplying by a matrix M as described in the main text amounts to a rescaling of the design matrix (in this case, B).

degree is strictly less than $|\Theta|$. This is clear from examining the right hand side of the equation above. While every term on the diagonal in this matrix is of the order $\lambda^{|\Theta|/|S|}$, every term *off* the diagonal is of the order $\lambda^{|\Theta|/|S|-1}$. To show that there is a term in the polynomial defined by the determinant that is of order less than $|\Theta|$, it suffices to show that some term in this expression reflects off diagonal terms. Note that the determinant is a sum over permutations $\sigma : |S| \rightarrow |S|$:

$$\sum_{\sigma} \text{sgn}(\sigma) \prod_{i=1}^n b_{i, b_{i, \sigma(i)}},$$

where $b_{i, \cdot}$ is the i th row of B . Then the permutations which simply flip one element (of which there are $|S| \cdot (|S| - 1)/2$ of) contribute to the determinant; since exactly two off diagonal terms contribute to this sum, the exponent on λ reflecting these permutations is distinct and less than $|\Theta|$.¹⁵

Therefore, as $\lambda \rightarrow 0$, we have that $(B^T B + \lambda I) \cdot \frac{1}{\lambda} I$ has a determinant that does not approach 0. Therefore, $(B^T B + \lambda I)^{-1} \lambda I$ has a limit (since in the definition of the matrix inverse, each term is scaled by the inverse of the determinant, and otherwise comes from multiplying and adding matrix elements together—so, since each term is scaled by a term that does not approach infinity, each term converges to a finite limit). Using Equation 1, we conclude that the limit defining $\tilde{\mathcal{I}}$ exists.

Now, note that the ridge estimator defined by $\tilde{\mathcal{I}}$ solves the following minimization problem:

$$\tilde{\mathcal{I}}_{\lambda}(\cdot)[s] = \underset{x}{\text{argmin}} \quad \|q_{s, \cdot} - Bx\|^2 + \lambda \|x\|^2. \quad (3)$$

By contrast, the information structure \mathcal{I} solves $q_{s, \cdot} = B\mathcal{I}(\cdot)[s]$. Now, as we have shown $\lim_{\lambda \rightarrow 0} \tilde{\mathcal{I}}_{\lambda}(\cdot)[s]$ exists; it follows from this expression that the resulting limit must also be a solution to the equation $q_{s, \cdot} = Bx$; if it weren't, then we would have the objective in (3) would converge to some strictly positive amount as $\lambda \rightarrow 0$; by contrast, *any* solution to this equation makes this objective equal to 0 in the limit. Hence any vector x which does not satisfy $q_{s, \cdot} = Bx$ cannot be the limit of $\tilde{\mathcal{I}}_{\lambda}(\cdot)[s]$ as $\lambda \rightarrow 0$.

On the other hand, for *any* vector x satisfying $q_{s, \cdot} = Bx$, subtracting the equation for $\tilde{\mathcal{I}}$ from this equation yields $0 = B(x - \mathcal{I}(\cdot)[s])$, so that $x - \mathcal{I}(\cdot)[s]$ is in the nullspace of B . But the information structure generating the decisionmaker's information is one possible choice of x ; therefore, $\mathcal{I}(\cdot)[s] - \tilde{\mathcal{I}}(\cdot)[s] = \sum_i \alpha_i v_i$, where $\{v_1, \dots, v_k\}$ is a basis for the nullspace of B (assuming the dimension of this space is k); adding $\tilde{\mathcal{I}}(\cdot)[s]$ to both sides of this expression proves the second half of the theorem. \square

Proof of the claim that an information structure is deterministic if and only if Q is the identity. A deterministic information structure involves the decisionmaker observing an element of the parti-

¹⁵More generally, the lowest degree of the polynomial should be $|\Theta| - |S|$; showing this, however, requires that some permutations which influence the determinant do not fix any elements on the diagonal. Determining that not all terms cancel out, while certainly intuitive, appears less direct than this argument. However, provided this is the case, then any entry corresponding to exclusively off-diagonal term will be a polynomial of order $(\lambda^{(|\Theta|-|S|)/|S|})^{|S|}$, since the matrix is $|S|$ -by- $|S|$.

tion of Θ . Suppose an information structure is partitional. This implies that the probability of observing any signal given any state is either 0 or 1. On the other hand, $b_{s,\theta}$ is positive if and only if $\mathcal{I}(\theta)[s] = 1$, meaning that $q_{s,\tilde{s}}$ is equal to 1 if $s = \tilde{s}$ and 0 otherwise. Therefore, Q is the identity.

Now suppose Q is the identity. Notice that each entry of Q is a convex combination of $\mathcal{I}(\theta)[\cdot]$, weighted according to a row of B . So, if $b_{s,\theta} > 0$, then we must have $\mathcal{I}(\theta)[s] = 1$. Notice that this immediately implies that this partitions the state space, since we cannot have two signals s, s' for which $b_{s,\theta} > 0$, since this would imply the rows of \mathcal{I} sum to a number greater than 1. Therefore, we obtain a partition of a subset of the state space Θ ; for any $\theta \in \Theta$ that is not in this subset, we have $b_{s,\theta} = 0$ for all $s \in S$. In this case, B and Q are generated according to a partitional information structure, where each element of the partition is the support of $b_{s,\theta}$ for some s , and where the prior assigns probability 0 to any state where $b_{s,\theta} = 0$ for all s . Hence, the information is generated by a partition. \square