The Impact of Network Structures on Information Processing*

Junya Murayama† Brian W. Rogers‡ Jacob VanNattan§
Xiannong Zhang¶

January 11, 2026

Abstract

We design experiments to study information processing through interconnected channels. About half of our subjects successfully back out the underlying information in a class of problems previously studied in the correlation neglect literature. We attribute their success to the application of a simple algorithm. Outside this class of problems, when the algorithm cannot be applied, yet Bayesian updating is similar, no subject solves any problem exactly, and those who originally succeeded fair no better than others. We also directly elicit comparative statics of a subject's estimate. Most subjects, including those who succeed in the baseline problems, fail to consistently identify the directional effects of signal changes on their solution, suggesting they do not understand the properties of the algorithm they use. We conclude that in problems that more closely resemble estimation tasks outside the lab, our subjects largely fail to understand some basic properties of Bayesian updating.

*We gratefully acknowledge funding from the Weidenbaum Center on the Economy, Government, and Public Policy at Washington University in St. Louis. We would like to thank Zachary Bell for performing excellent data analysis for this project.

†Washington University in St. Louis, U.S.A., j.murayama@wustl.edu
‡Washington University in St. Louis, U.S.A., brogers@wustl.edu
§Washington University in St. Louis, U.S.A., j.d.vannattan@wustl.edu
¶Washington University in St. Louis, U.S.A., zxn@google.com

1 Introduction

We are interested in situations where a decision maker accesses information through several channels and must consider how to combine the resulting signals. Before a consumer makes a purchase, they may observe professional and user reviews evaluating various options. Similarly, an investor can acquire ratings and sector forecasts before taking a financial position. Opinions on social phenomena such as vaccine efficacy, political policies, and historical events are often influenced by views expressed on social and traditional media platforms. The signals arising through such channels are generally interdependent, as various channels rely partially on common sources.

We design and implement a battery of experiments in an environment meant to capture an abstract representation of these examples. The decision problems are based on a simple concrete mechanism in which the signal structure is modeled through an exogenous network. The subject's task is to estimate an average of several (normally distributed) i.i.d. information sources. Rather than directly observing the sources, the subject observes a set of signals, where each signal produces an average of a given subset of sources. In this context, the optimal estimate is a linear function of the signals, though generally not a simple average, as signals can (i) have different precisions due to averaging different numbers of sources, and (ii) be positively correlated due to sharing sources in common.

At the heart of our design are four structures that we call "Conclusive Principal Structures", which include the main structure from Enke and Zimmermann (2019). Each structure can be represented as a matrix that encodes which signals are connected to which sources and can be solved through elementary linear methods. But from a behavioral perspective, these problems are non-trivial for many people. About half (49%) of our observations in the Conclusive Principal Structures are rational. A preliminary but crucial conclusion is that the remaining, non-rational responses are highly heterogeneous and inconsistent across tasks at the individual subject level. Previous literature has focused on the role of correlation neglect in driving a wedge between Bayesian updating and observed behavior and concluded that an extreme form of such a bias can explain much about non-rational behavior. Our data is similar across the four Conclusive Principal Structures, and the hypothesis that behavior can be largely explained by correlation neglect, or by a simple averaging of the signals, is strongly rejected.

We identify a second bias, which we call precision neglect, in which a decision-maker under-appreciates the differential informativeness of signals. In two sequences of the simplest possible problems that allow us to explicitly elicit a subject's understanding of precisions and correlations, we find that both correlation neglect and precision neglect are important determinants of behavior. Yet even the combination of these two biases is not nearly enough to provide a convincing account of the distribution of responses in the Conclusive Principal Structures. Furthermore, since our design is within-subject, we are able to estimate the joint distribution of the two biases in our population, and show that these biases are strongly related to behavior in the Conclusive Principal Structures.

These findings lead us to novel and qualitatively different conclusions about how experimental subjects process related channels of information. The remainder of our analysis is built around the theme of understanding more deeply *how* subjects think about updating when confronted with interdependent sources of information. Our central finding is that subjects' ability to solve these problems relies on implementing a simple algorithm that applies only when posterior beliefs are degenerate. As such, our subjects' success is remarkably fragile, and does not carry over to settings that are similar, yet require a deeper understanding to solve. We conclude that their partial success in the Conclusive Principal Structures lacks robustness and is unlikely to meaningfully extend into areas of normal life.

The findings that support this conclusion are based on two extensions of the experimental design beyond the Conclusive Principal Structures. The first extension is to a set of four Inconclusive Principal Structures, each of which is obtained by removing one signal from its paired Conclusive Principal Structure. In these problems, the Bayesian posterior is non-degenerate, i.e., the realizations of the sources cannot be exactly determined from the signals. The eight Principal Structures are depicted in Figure 2.

The estimates subjects provide in the Inconclusive Structures are quite different. None of our subjects exactly solve these problems, despite the fact that the Bayesian posterior expectation remains a linear calculation. Furthermore, the subjects who succeed in the Conclusive Principal Structures fair no better in the Inconclusive Principal Structures, on average, than the others. The Conclusive Structures can be solved via a simple algorithm that identifies a source whose realization can be readily and exactly determined. This value can then be used in combination with another signal to exactly determine the value of a second source, and so on, until the realizations of all sources are known, after which they are averaged. For example, in CON1 of Figure 2, Source 2 can be directly inferred from Signal 2. Once Source 2 is known, one can use Signal 1 to deduce Source 1, and can use Signal 3 to deduce Source 3. We suggest the following explanation of our data. About half of our subjects identify and apply this algorithm, arriving at correct answers to Conclusive Structures. However, the ability to identify and implement this algorithm confers no ability to solve an Inconclusive Structure, nor even any marginal advantage over subjects who do not use the algorithm.

The final experimental domain is designed to identify how subjects think about the comparative statics of their estimates. Rather than ask a subject for a point estimate corresponding to a particular realization of signals, we instead effectively ask the subject to provide a linear function of the signals before they are realized. The subject provides the coefficients, which dictate the partial derivatives of their estimate with respect to each signal. This methodology addresses the identification problem whereby, in the baseline case, it is not obvious how each signal influences a subject's estimate. In this *ex ante* form, subjects directly report the weight of each signal in their final estimate.

We find that very few of the subjects understand the appropriate weights to attach to signals, which carries the implication that they do not understand the comparative statics properties of the problems they face. Not only do the assigned weights differ quantitatively from those of a Bayesian, but subjects often assign weights with the opposite sign. Although subjects who succeed in the Conclusive Principal Structures are more likely to answer these questions qualitatively correctly, between two-thirds and three-fourths of such subjects assign weights with the wrong signs. We conclude that, for most subjects, the ability to solve a specific instance of a problem does not signify an understanding of the basic qualitative features of the problem.

The success we observe in the Conclusive Principal Structures, which is very much in line with previous literature, thus seems to have fairly narrow limits. Our view is that the ability to apply the algorithm described above is not likely to be important in problems outside the lab where interdependent information sources must be combined to form a posterior. Instead, the skills likely to translate to more general success are related to having a sound qualitative understanding of the implications of correlations and precision differences across information channels. Possession of these skills is rare in our data and mostly not signified by success in the Conclusive Principal Structures. We suggest that the Inconclusive Structures and the *ex ante* form of elicitation each yield better proxies for measuring those skills, and we therefore reach much more conservative conclusions about subjects' understanding of these problems.

We are not the first to propose an experimental study in this kind of environment. The most closely related work is Enke and Zimmermann (2019) (Hereafter

EZ), which is largely responsible for inspiring this work, and from which we borrow certain important features of the experimental environment. The robust finding from EZ is that approximately half of their subjects are rational, while the other half display a bias in the direction of only partially accounting for the correlations across reports, which they call "correlation neglect". Our baseline experimental design, and methods of data analysis, offer several important improvements over the approach in EZ and are responsible, in part, for our substantively different conclusions. One important observation is that, in EZ's analysis, precision and correlation biases are implicitly bundled together, and so there is no attempt to disentangle the two phenomena.

Beyond the immediate connection with EZ, our paper contributes to the literature studying systematic errors in belief formation (see e.g. Moore and Healy (2008); Esponda and Vespa (2014); Charness and Levin (2009); Charness *et al.* (2010); Hanna *et al.* (2012); Garfagnini and Walker-Jones (2024); Agranov and Reshidi (2024)). While this is a large and nuanced literature, one of its broad goals is to understand how different biases contribute to behavioral phenomena that arise from suboptimal information processing. We contribute to this agenda by separately identifying a bias that relates to differential signal quality, "precision neglect", and one that fails to account for common sources driving different signals, "correlation neglect". We jointly identify their distribution among our subjects, showing how the biases combine to determine behavior, and how these conclusions are affected by the statistical properties of the underlying problem. In terms of the classification in the literature review provided in Enke (2024), we find evidence that subjects use noisy approximations and that, in certain cases, they resort to solving a simpler but related problem.

The behavioral patterns we identify in our experiments can have an important impact on overconfidence. Ortoleva and Snowberg (2015) point out that neglecting correlations among signals leads to overconfidence in one's own beliefs, which can translate into ideologically extreme views. Moore and Healy (2008) find that the inability to fully understand signal precision is an important aspect of overconfidence, which they call overprecision. **?** find that subjects overweight redundant information. The same phenomenon occurs in certain versions of our treatments. In particular, we find that many subjects place positive weight on redundant information (see Section 4.3).

We also contribute to the recent literature that studies behavioral responses to complexity in the decision-making environment. For example, Arrieta and Nielsen (2024) show that as a problem becomes more complex, the decision-making process becomes easier to describe. While our notion of complexity is different, we also find evidence that as problems become more complex, as in the distinction between our Inconclusive and Conclusive problems, many subjects resort to simpler strategies.

Our setting can be interpreted along the lines of social learning in networks (see e.g. Choi *et al.* (2012); Chandrasekhar *et al.* (2020); Grimm and Mengel (2020); Eyster *et al.* (2015); Eyster and Rabin (2014, 2010); DeMarzo *et al.* (2003)). However, note that our setting is simpler than a classical social learning environment for at least two reasons. First, our subjects receive signals that are simple averages of the underlying sources, rather than other agents' actions, so the belief formation process does not involve higher-order beliefs. Second, our subjects only answer one question in each network structure, so the environment is static rather than dynamic. Our paper focuses on how agents process information from a fixed network, which can be seen as a fundamental part of dynamic social learning problems that involve multiple agents and dynamic effects. In this simplified setting, we find empirical evidence that is very much consistent with this literature. For instance, Chandrasekhar *et al.* (2020) find that around 50% of subjects are a "Bayesian type". We similarly find about half of observations are rational in our Conclusive Principal Structures. Grimm and Mengel (2020) point out that both Bayesian and DeGroot learning models have limited power in predicting behavior in their setting and find that subjects take the correlation of neighbors into account in a more rudimentary way than a Bayesian would. Their finding is consistent with our results, in which agents are heterogeneous in their understanding of signal precision and correlation. Our environment is also related to the persuasion bias proposed by DeMarzo *et al.* (2003). We find empirical evidence consistent with this effect, where agents fail to account for repetition of information, leading to well-connected sources having an exaggerated impact on beliefs.

Still, the most closely related work to our own focuses on behavior directly related to correlation neglect. Eyster and Weizsacker (2016) study the consequences of correlation neglect for financial decision-making. They directly test subjects' attention to correlation and find subjects tend to ignore correlations and follow a simple "1/N heuristic". We find that this heuristic is followed by some subjects in more complex problems. We clarify that this simple averaging does not arise purely through correlation neglect, but instead arises from the combination of correlation neglect and precision neglect. Hossain and Okui (2024) study overprecision and correlation neglect in financial decisions. Similar to our point, they argue that correlation neglect should not be studied in isolation from misperception of signal variance. They find that a model incorporating over-precision fits their data best. In contrast, we do not find systematic over- or under-estimation of precision when we estimate the distribution of biases associated with precision. Instead, we find that subjects are inattentive to *differences* in precision across information from different channels.

The rest of the paper is organized as follows. Section 2 summarizes the relevant theory. Section 3 describes the experimental design. Section 4 presents the experimental results. Section 5 concludes.

2 The decision-making environment

We now formally define the problems our subjects face. Let $\boldsymbol{x} \sim N(\mathbf{0}, \boldsymbol{\sigma^2 I_n})$ be a vector of i.i.d normal variables. We call a realization of $\boldsymbol{x}$, $\hat{\boldsymbol{x}}$, source information, or just *sources*. Let $\hat{\boldsymbol{y}} = \boldsymbol{\Gamma}\hat{\boldsymbol{x}}$ be a linear transformation of $\hat{\boldsymbol{x}}$, which we call *signals*. We call the $m \times n$ matrix $\boldsymbol{\Gamma}$ an *information network*. We assume throughout that $n \geq m$, that the sum of each row and column of $\boldsymbol{\Gamma}$ is strictly positive, and that $\boldsymbol{\Gamma}$ has full row rank. Let $\boldsymbol{t} = \frac{1}{n}\mathbf{1}$ be the $n \times 1$ *target vector*, so that $\boldsymbol{t'}\hat{\boldsymbol{x}}$ represents the mean of the sources $\frac{1}{n}\sum_{i=1}^{n} \hat{x}_i$.

A decision maker who knows $\boldsymbol{\Gamma}$ and $\boldsymbol{x} \sim N(\mathbf{0}, \boldsymbol{\sigma^2 I_n})$ is incentivized (through a quadratic loss payoff function) to estimate $\boldsymbol{t'}\hat{\boldsymbol{x}}$ given $\hat{\boldsymbol{y}}$. Formally, the subject's problem (P) is

$$\arg\max_{a \in \mathbb{R}} \mathbb{E}[\pi(a, \boldsymbol{t'}\hat{\boldsymbol{x}})] = \int \pi(a, \boldsymbol{t'}\hat{\boldsymbol{x}}) L(\hat{\boldsymbol{x}}|\hat{\boldsymbol{y}}) d\hat{\boldsymbol{x}} \tag{P}$$

where $\pi(a, \boldsymbol{t'}\hat{\boldsymbol{x}}) = -(a - \boldsymbol{t'}\hat{\boldsymbol{x}})^2$, and $L(\hat{\boldsymbol{x}}|\hat{\boldsymbol{y}})$ is the conditional likelihood of $\hat{\boldsymbol{x}}$ given $\hat{\boldsymbol{y}}$ under the prior information. Let $a^* = \arg\max_{a \in \mathbb{R}} \mathbb{E}[\pi(a, \boldsymbol{t'}\hat{\boldsymbol{x}})]$ denote the solution to this problem.

Proposition 1. *We relate the subject's estimation problem to the following distance-minimization problem:*

$$\min_{\hat{\boldsymbol{x}} \in \mathbb{R}^n} \hat{\boldsymbol{x}}'\hat{\boldsymbol{x}} \tag{P1}$$

$$s.t.\ \boldsymbol{\Gamma}\hat{\boldsymbol{x}} = \hat{\boldsymbol{y}}$$

The solution to (P1) is given by

$$\hat{\boldsymbol{x}}^* = \boldsymbol{\Gamma'}(\boldsymbol{\Gamma\Gamma'})^{-1}\hat{\boldsymbol{y}} \tag{1}$$

and the solution to (P), the optimal a, is given by

$$a^* = \boldsymbol{t'}\boldsymbol{\hat{x}^*} \tag{2}$$

Proof. See the appendix for a proof of a more general result. □

By assumption $\boldsymbol{\Gamma}$ has full row rank and so $\boldsymbol{\Gamma\Gamma'}$ is invertible. Thus for any signals $\boldsymbol{\hat{y}}$ and information network $\boldsymbol{\Gamma}$ that satisfy our assumptions above, Proposition 1 transforms a normal-quadratic utility maximization problem into a simpler distance minimization problem, showing that the Bayesian estimate identifies the $\boldsymbol{\hat{x}^*}$ that is closest to the prior, among the possibilities consistent with signals $\boldsymbol{\hat{y}}$. The solution $\boldsymbol{\hat{x}^*}$ is not necessarily equal to the realized sources $\boldsymbol{\hat{x}}$, and, in fact, in some cases the decision maker is not able to recover $\boldsymbol{\hat{x}}$ from the available information. Nevertheless the Bayes rational estimate $a^* = \boldsymbol{t'}\boldsymbol{\hat{x}^*}$ is of course well-defined. Note also that $\boldsymbol{\hat{x}^*}$ is independent of σ^2, the prior variance.[1]

We call a pair $(\boldsymbol{\Gamma}, \boldsymbol{t})$ *conclusive* if there exists a vector $\boldsymbol{\tilde{t}}$ $(m \times 1)$ such that $\boldsymbol{t'} = \boldsymbol{\tilde{t}}'\boldsymbol{\Gamma}$. In this case $\boldsymbol{t'}\boldsymbol{\hat{x}} = \boldsymbol{\tilde{t}}'\boldsymbol{\Gamma}\boldsymbol{\hat{x}} = \boldsymbol{\tilde{t}}'\boldsymbol{\hat{y}}$, so $a^* = \boldsymbol{\tilde{t}}'\boldsymbol{\hat{y}}$ recovers the mean of the sources exactly (without necessarily determining the value of each source). Note that conclusivity is weaker than invertibility of $\boldsymbol{\Gamma}$. If $\boldsymbol{\Gamma}$ is invertible, then we simply have $\boldsymbol{\hat{x}} = \boldsymbol{\Gamma^{-1}}\boldsymbol{\hat{y}}$, so that the vector x can be exactly determined from y. All other $(\boldsymbol{\Gamma}, \boldsymbol{t})$ pairs are called *inconclusive*. An important distinction is that, for inconclusive problems, the posterior beliefs are non-degenerate.

In addition to Bayes rational estimates, we will consider a naive model in which the signals are simply averaged, effectively ignoring $\boldsymbol{\Gamma}$. We call this behavior *Uniform*. Notice that such behavior disregards both the correlations between the elements of $\boldsymbol{\hat{y}}$ as well as the fact that they generally have different variances. For this reason, we distinguish between *correlation neglect* and a separate though potentially related bias, which we term *precision neglect*. See the appendix for formal definitions of these behavioral patterns.

We now define notions of precision and correlation for a given information network $\boldsymbol{\Gamma}$. We use *connectivity*, κ_i, to describe the precision or informativeness of signal $\hat{y}_i$, $i \in \{1, \dots, m\}$.

$$\kappa_i = \sum_{j=1}^{n} \mathbb{1}_{\{\boldsymbol{\Gamma}_{ij} > 0\}} \tag{3}$$

Graphically, κ_i is simply the number of sources that signal i is connected to.

[1] Thus base-rate neglect, modeled as a decision-maker who believes $\boldsymbol{x} \sim N(\boldsymbol{0}, s^2\boldsymbol{I_n})$, where $s^2 > \sigma^2$, does not affect behavior, since $\boldsymbol{\hat{x}^*}$ is independent of s^2.

We use *overlap*, $\omega_{ii'}$, to describe the correlation between signals $\hat{y}_i$ and $\hat{y}_{i'}$, $i, i' \in \{1, \ldots, m\}$, $i \neq i'$, as

$$\omega_{ii'} = \sum_{j=1}^{n} [\mathbb{1}_{\{\mathbf{\Gamma}_{ij}>0\}} \cdot \mathbb{1}_{\{\mathbf{\Gamma}_{i'j}>0\}}] \tag{4}$$

Graphically, $\omega_{ii'}$ is simply the number of sources that signals i and i' share in common.

3 Experimental design

The design has two treatments, which we refer to as Precision-Correlation (PRE-COR) and Ex Ante. Both treatments share in common the eight Principal Structures. The PRE-COR Treatment includes two additional sequences of problems, which we refer to collectively as the PRE-COR Structures. The Ex Ante Treatment presents the Principal Structures in two additional phases after the initial phase. The second phase presents the structures in an *ex ante* format, where we elicit a weight for each signal separately. The third phase assesses subjects' confidence in their original estimate compared to the implied estimate from the *ex ante* phase. See Figure 1 for a summary.

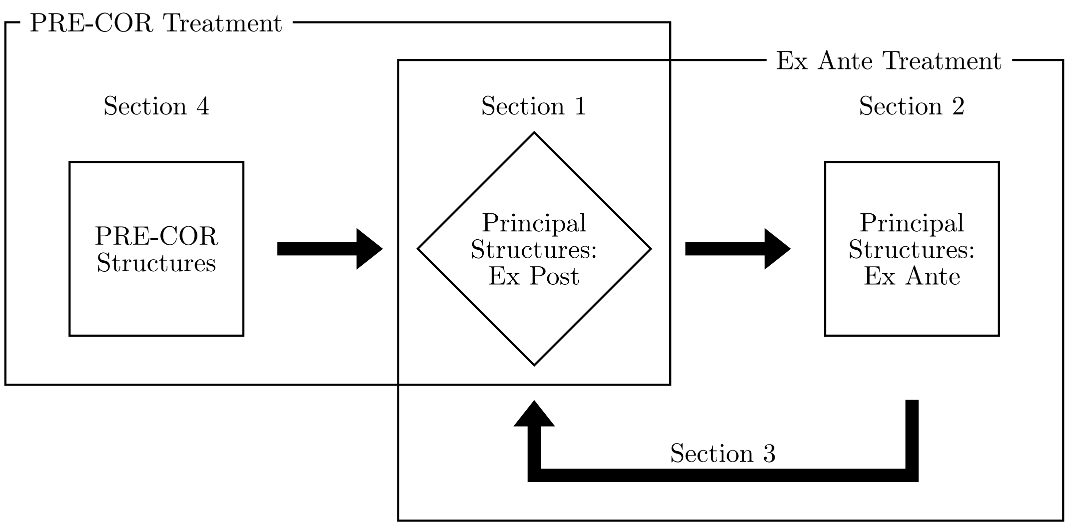

Figure 1: Summary of the experimental design.

The experiments were conducted asynchronously via Qualtrics, through which we administered all instructions, a questionnaire, and payment. 99 subjects were recruited through the Missouri Social Science Experimental Laboratory Subject Pool. In each section listed in Figure 1, a random question is selected at the end of the experiment and subjects receive $5 - 0.2(a - a^*)^2$ dollars, where a^* is the rational solution and a is a subject's answer. In the Ex Ante Treatment, since we use the same realizations of Principal Structures in each section, we randomly select 3 *different* structures across sections. Each session lasted about 60 minutes, and the average payment was 14.75 USD.[2]

3.1 Principal structures

At the core of our design are four *Conclusive Principal Structures* and a set of four paired *Inconclusive Principal Structures*. Subjects are given a graphical representation of the information network as in Figure 2, and numerical signal realizations $\hat{y}$ (the circles) that are linear transformations (in fact, simple averages) of a subset of sources $\hat{x}$ (the squares) indicated by the links in the diagram. They are told, in simple terms, that the sources are i.i.d. draws from a standard normal distribution. They are asked to estimate the average of all sources, with payment dictated by quadratic loss between the true value $\bar{x}$ and their estimate, with a lower bound at zero. Note that while our CON3 uses the same network as the main structure employed by EZ, there are three important differences between their design and our Principal Structures. First, our subjects are incentivized to estimate the sample mean $\boldsymbol{t'\hat{x}}$ from a known distribution as opposed to estimating an unknown population mean. Second, because our underlying distribution is fixed and known to the decision-maker, we avoid the inference complications associated with tying the mean and variance of the underlying, unknown, distribution together. Third, our payment rule is quadratic loss, as opposed to quadratic proportional loss, so that subjects are incentivized to report the mean of their posterior beliefs. These differences bring our experimental design much closer to EZ's "Robustness Treatment".

Of the four Principal Structures listed in the top row of Figure 2 (CON1 - CON4), only CON2 has a matrix that is not invertible. Each Conclusive Principal Structure is paired with an Inconclusive counterpart (INC1 - INC4) by removing one link and one signal; these are depicted in the bottom row of Figure 2. Proposition 1 shows that the Bayes-rational estimate is a linear combination of the signals and we call the coefficients attached to the signals *weights*. Notice

[2]Complete screenshots from the both treatments are available at https://wustl.box.com/s/svqgzi0v6rgbkhqqsm3gsgi0nl217kuf.

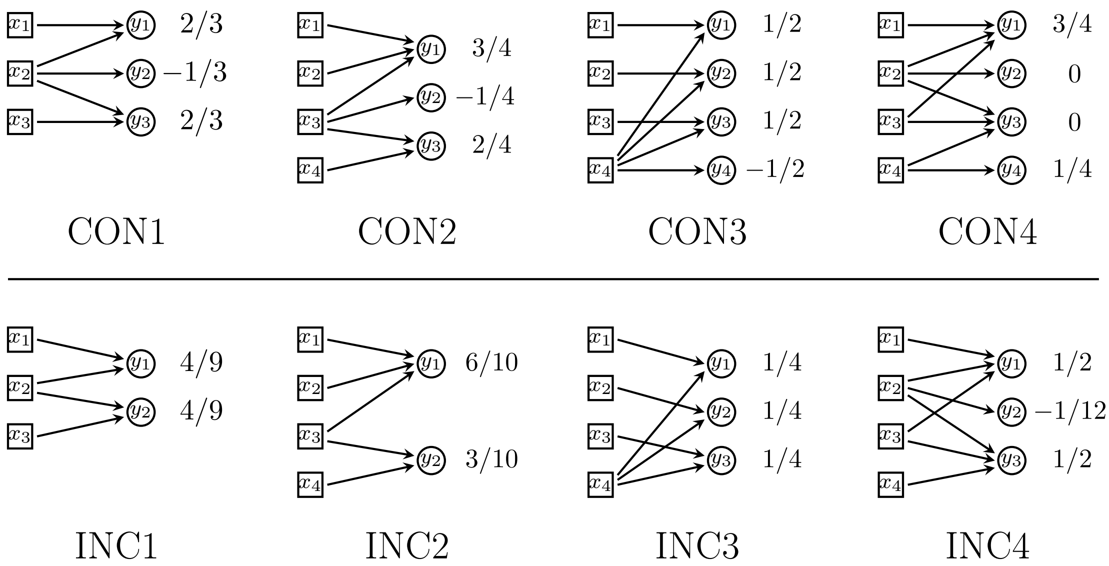

Figure 2: The four pairs of Conclusive and Inconclusive Principal Structures along with the Bayes-rational weights attached to the signals.

that the number of signals and links contained in a Principal Structure strictly increases from the left column of Figure 2 to the right column.

We now summarize the solutions to, and main properties of, the Principal Structures.

Pair 1. The first pair contains the smallest conclusive and inconclusive problems. In fact, INC1 is the simplest possible inconclusive structure.[3].

For CON1, we have $\mathbf{\Gamma} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 \\ 0 & 1 & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}$, $\boldsymbol{t} = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}'$.

Since $\mathbf{\Gamma}$ is invertible, the solution can be obtained from $\hat{\boldsymbol{x}} = \mathbf{\Gamma^{-1}}\hat{\boldsymbol{y}}$. But notice that the following algorithm can be applied. A subject can start by noticing that y_2 directly reports x_2, i.e., $x_2 = y_2$. Once x_2 is known, y_1 and x_2 can be used to compute $x_1 = 2y_2 - x_2$, and then, similarly, y_3 and x_2 can be used to compute $x_3 = 2y_3 - x_2$. We summarize this algorithm with the notation $y_2 \to x_2 \to y_1 \to x_1 \to y_3 \to x_3$. Notice that no such algorithm exists for INC1. Still, INC1 can be represented as $\mathbf{\Gamma} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}$, $\boldsymbol{t} = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}'$, and

[3]This is true under the following conditions: (1) All sources are connected to at least one signal. (2) All signals are connected to at least one source. Then, a network needs to contain at least three sources and two signals to be Inconclusive. Therefore INC1 is the simplest Inconclusive structure in the sense it contains the least possible number of signals, sources, and links.

Proposition 1 can be directly applied to compute the rational weights shown in Figure 2. When we turn to the data, we note that it is possible, in principle, to apply the algorithm described above, without solving CON1 in the sense of Proposition 1 to obtain the rational weights depicted in Figure 2. Notice also that the sum of the rational weights in any conclusive problem is 1, while the sum of the weights in any inconclusive problem is strictly less than one, which is manifested in this pair of problems, as well as in the following pairs.

Pair 2. CON2 is the only conclusive structure in our problems that is not invertible. As a consequence, notice that while $(x_1 + x_2)/2$ can be exactly determined, and hence the target $(x_1 + x_2 + x_3 + x_4)/4$ can be exactly determined, it is not possible to recover the exact realizations of either x_1 or x_2. For CON2, we have $\boldsymbol{\Gamma} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}$, $\boldsymbol{t} = \begin{pmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}'$.

The algorithm that can be used to solve CON2 in place of Proposition 1 is $y_2 \to x_3 \to y_1 \to (x_1 + x_2) \to y_3 \to x_4$.

INC2 is represented as $\boldsymbol{\Gamma} = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}$, $\boldsymbol{t} = \begin{pmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}'$.

Pair 3. CON3 is the same structure employed by EZ. It can be solved by $y_4 \to x_4 \to y_3 \to x_3 \to y_2 \to x_2 \to y_1 \to x_1$. Alternatively, we can write

$$\boldsymbol{\Gamma} = \begin{pmatrix} \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & 1 \end{pmatrix}, \boldsymbol{t} = \begin{pmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}'.$$

INC3 is obtained by removing signal y_4. We have

$$\boldsymbol{\Gamma} = \begin{pmatrix} \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}, \boldsymbol{t} = \begin{pmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}'.$$

Note that, as with INC1, the signals in INC3 are symmetric in structure, and so the rational weights are equal across signals.

Pair 4. The final pair of problems are the largest, in terms of number of links and signals. CON4 is the only principal structure to encode redundant information. Notice that the optimal weights on signals 2 and 3 is zero. Indeed, signal 1 reports the average of the first three sources and signal 4 reports the fourth source, from which the optimal estimate is readily obtained. The algorithm is $y_4 \to x_4 \to y_1 \to (x_1 + x_2 + x_3)$. The presence of signals 2 and 3, while not necessary to recover the target average, do make the structure invertible, as opposed to merely conclusive, since they allow an exact determination of each source. Thus, an alternative algorithm is $y_2 \to x_2 \to y_4 \to x_4 \to y_3 \to x_3 \to y_1 \to x_1$. The representation of CON4 is

$$\boldsymbol{\Gamma} = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & 1 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & 0 & 0 & 1 \end{pmatrix}, \boldsymbol{t} = \begin{pmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}'.$$

INC4 is obtained by removing signal 4. Notice that INC4 is the only inconclusive structure with a negative weight. It is formally represented as

$$\boldsymbol{\Gamma} = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & 1 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}, \boldsymbol{t} = \begin{pmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}'.$$

3.2 The Ex Ante Treatment

The Ex Ante treatment consists of three sections, numbered Sections 1-3 in Figure 1, each of which uses the eight Principal Structures described above. In Section 1, the problems are presented as described: the subject is given the structure of the problem, a realization of signals, and is asked to provide an estimate of the target average. We refer to this as the *ex post* form. The *ex post* form is the baseline; all Principal Structures are asked in this form at least once to all subjects in both treatments. In Section 2, subjects are presented with the same eight structures, but in *ex ante* form. The important distinction is that subjects are asked to estimate the target average before the signals are realized. Thus subjects are asked to provide a weight for each signal, according to which their estimate will be constructed as the corresponding linear combination of signals once they are realized. The elicitation mechanism is a graphical slider bar, with a range of $[-1, 1]$ and with no initial value selected as a default. The subject also has a slider bar to attach to zero, the mean of the prior distribution. In total, there are m slider bars for signals and 1 slider bar for the prior, so that $m + 1$ quantities are elicited for each Principal Structure. Before answering *ex ante* questions, subjects are told, in simple terms, that the optimal solution to these questions can be achieved by linearly combining the signals. In Section 3, we repeat the same problems in *ex post* form using the same signal realizations as Section 1. In these repeated problems, subjects are reminded of their estimates in Section 1 as well as their implied estimates using the weights they provided in Section 2. Subjects are free to provide any estimate and are not constrained to choose one of their previous answers. All feedback and payment information is provided at the end of the experiment to minimize learning effects.

3.3 The PRE-COR Treatment

The PRE-COR Treatment also contains the eight Principal Structures (in *ex post* form). However, it begins with two sets of estimations, the Precision Sequence (PRE) and the Correlation Sequence (COR), both of which are depicted in Figure 3.

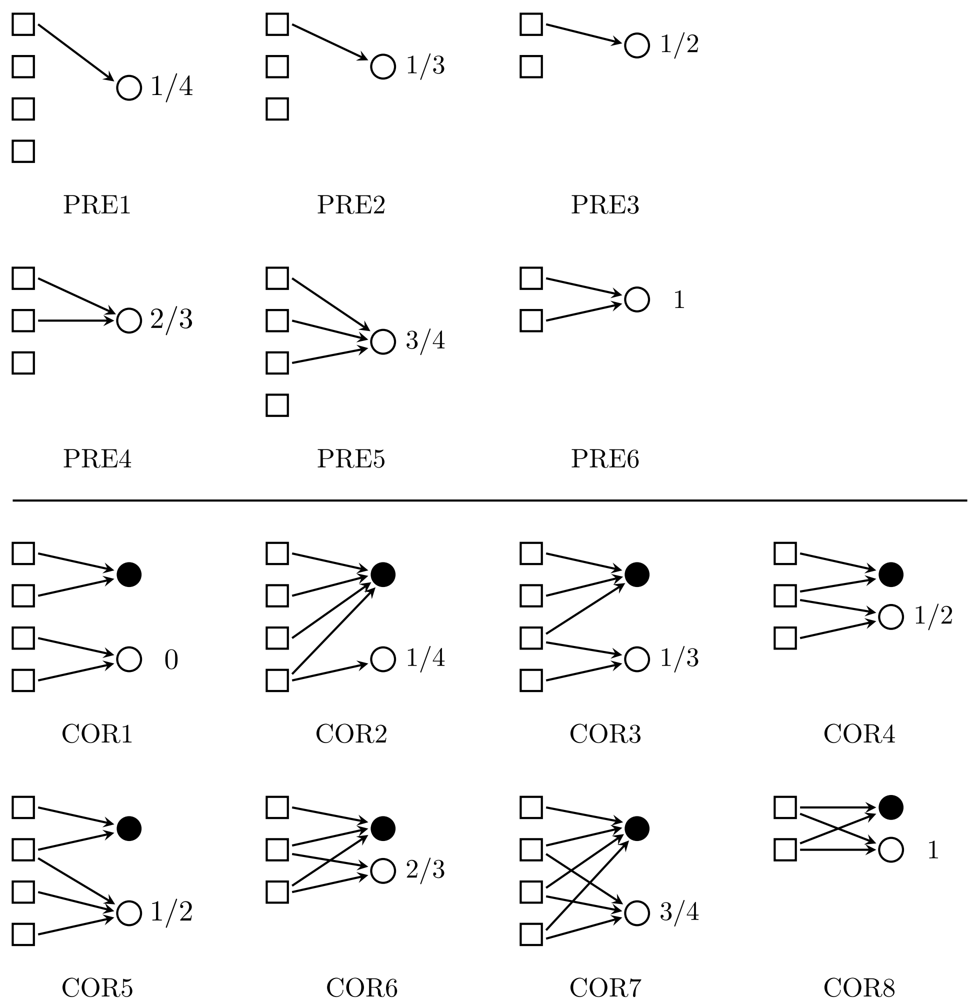

Figure 3: PRE-COR Structures

The Precision Sequence is designed to measure, in a simple way that is closely related to the Principal Structures, how well a subject understands variation in the variance of a signal. Subjects are asked to estimate the average number in a set of sources (the squares) given the value of a single signal (the circles). The rational weight to attach to the signal i, $\frac{\kappa_i}{number\ of\ sources}$, strictly increases from PRE1 to PRE6.

The second set of questions, the Correlation Sequence (COR), is depicted in the bottom two rows of Figure 3. Subjects are asked to estimate the *value of a signal* (the black circles) given the value of the other signal (the white circles). As before, the rational weight, $\frac{\omega_{ii'}}{\kappa_{i'}}$, weakly increases from COR1 to COR8, with strict increases except for COR4 and COR5.

The design of the PRE-COR Treatment is summarized in the left-hand-side of Figure 1. As in the Ex Ante Treatment, feedback and payoff information is not revealed until the end of experiment to minimize learning effects.

4 Results

We begin with results from the four Conclusive Principal Structures at the core of our design in Section 4.1. We then demonstrate, in Section 4.2, that behavior in the Inconclusive Principal Structures is dramatically different, and we relate the behavior in the Principal Structures to measures of understanding derived from the PRE-COR Structures. We report the results of the Ex Ante Structures, where behavior is again dramatically different from the CPS data, in Section 4.3.

4.1 Conclusive Principal Structures

Nearly one-half of subjects' estimates in the four CPS are Bayes-rational: each of 99 subjects faces four CPS problems and 180 out of 396 observations are Bayes-rational. CON3 is the main structure employed by Enke and Zimmermann (2019), and our finding that about half of subjects succeed is very similar to theirs, as well as to Chandrasekhar *et al.* (2020). Moreover, this proportion is not special to this particular structure: it is remarkably stable across the four problems. The success rate in CON1, CON3, and CON4 range from about 45% to about 48%. The success rate in CON2 is marginally lower, at about 40%, and this makes sense. CON2 is the only CPS in which it is not possible to individually recover all source realizations. Specifically, while the mean of sources 1 and 2 can be recovered exactly, the individual realizations of those two sources cannot. We conjecture that this distinction prevented some subjects from solving the problem. At the individual level, slightly more than half of our subjects (53/99) provide two or more Bayes-rational estimates, and about one-fourth (24/99) correctly solve all four problems. The Wilson Score Test gives 95% confidence intervals on these figures of [0.44%,0.63%] and [0.17%,0.34%] respectively.

Turning to the complementary observations, where subjects do not report the

Bayes-rational estimate, our main conclusion is that this data is highly idiosyncratic and cannot be readily explained by any simple model. The naive model of simply averaging the reports, which we call Uniform, accounts for less than 10% of the data in CPS.[4]

Structure	BR (0.25 / 0.05)	U (0.25 / 0.05)
CON1	48.5% / 45.5%	16.2% / 12.1%
CON2	45.5% / 40.4%	6.1% / 5.1%
CON3	53.5% / 48.5%	14.1% / 10.1%
CON4	47.5% / 47.5%	11.1% / 7.1%
Pooled	48.7% / 45.5%	11.9% / 8.6%

Table 1: Conclusive Principal Structure response frequencies

Table Notes: Data includes all 99 subjects from both the PRE-COR and Ex Ante treatments. This table presents the percentage of subjects whose estimates for each structure fell within a 0.25 and 0.05 radius of the point prediction of the Bayes-rational (BR) and Uniform (U) models.

We collect these observations in the following:

Result 1: Across the four CPS, consistently, nearly half of the responses are rational, while the complementary data are highly idiosyncratic.

4.2 Inconclusive Principal Structures

Recall that each CPS problem is paired with an IPS problem via the removal of a single report and link. In each CPS, the Bayes-rational estimate can be constructed through a simple algorithm that begins by using one report to exactly identify one source. Once the first source is known, a second report can be used to exactly identify a second source, and so on, until each source is known, after which the mean is easy to compute.

A crucial aspect of the distinction between Conclusive and Inconclusive problems is that this algorithm fails. Indeed, neither the individual sources nor the mean of these sources can be exactly identified in the IPS problems. Notice that this distinction is not necessarily obvious in the representation of the problem's solution given in Equation (1).

This observation raises the following empirical question. When subjects correctly estimate the mean of sources in the CPS problems, do they succeed because

[4]Based on our independent analysis, this finding is also consistent with the raw data from Enke and Zimmermann (2019).

they understand the content of Equation (1), or do they succeed through the application of the algorithm above? If it is the former, then the data in the IPS problems should look similar to that in the CPS problems. If it is the latter, then the data should reveal a different pattern of responses.

This distinction is fundamental to the interpretation of the CPS data. The algorithm that can be applied to the CPS problems is arguably artificial and specific to the highly stylized problems presented in the laboratory. Conversely, any application of these ideas outside the laboratory is likely to be more closely related to the IPS problems. Combining interdependent sources of information from several experts requires a qualitative understanding of how their reports should influence beliefs. In the absence of an explicit and precise statistical framework, and where signals are likely to be transmitted noisily, most real-world problems are inherently of the inconclusive variety.

Our central finding here is that subjects approach the IPS problems in a qualitatively different manner from the CPS problems. This difference has two main consequences. First, the estimates in the IPS problems are much less precise. Second, subjects are much closer to homogeneous in their ability to solve the IPS problems; there is no longer any clear way to separate them in terms of success.

More specifically, when success requires high precision, performance in the IPS problems is drastically lower than in the CPS problems. Out of the $4 \cdot 99 = 396$ observations, only 13 estimates (3.3%) are within 0.05 of the Bayes-rational report, whereas the corresponding figure from the CPS problems is $180/396 \approx 45.5\%$. No subject produced more than one approximately rational estimate.[5] We conclude that, in sharp contrast to behavior in the CPS problems, none of our subjects are exactly solving any of the IPS problems.

Apparently, the general solution represented by Equation (1) is elusive, and some of our subjects can find it only when the algorithm can be applied.[6] Subjects seem to be instead using intuition to produce their estimates. Figure 4 summarizes the distribution of responses in the Principal Structures. While their estimates are not exact, they are perhaps surprisingly well distributed around the Bayes-rational estimate. Notice also that the the naive behavior of averaging the signals, which we call Uniform, becomes the modal response in IPS problems. To rationalize Uniform averaging, a subject must treat the signals as i.i.d. or, in other words, to

[5]Using a wider interval of ± 0.25, there was one subject (out of 99) who produced 3 estimates in the vicinity of the rational estimate, perhaps demonstrating a strong intuition.

[6]Of course, our conclusions may or may not generalize to a setting where the incentives are much higher.

ignore both the correlations between them and the differences in their precision.

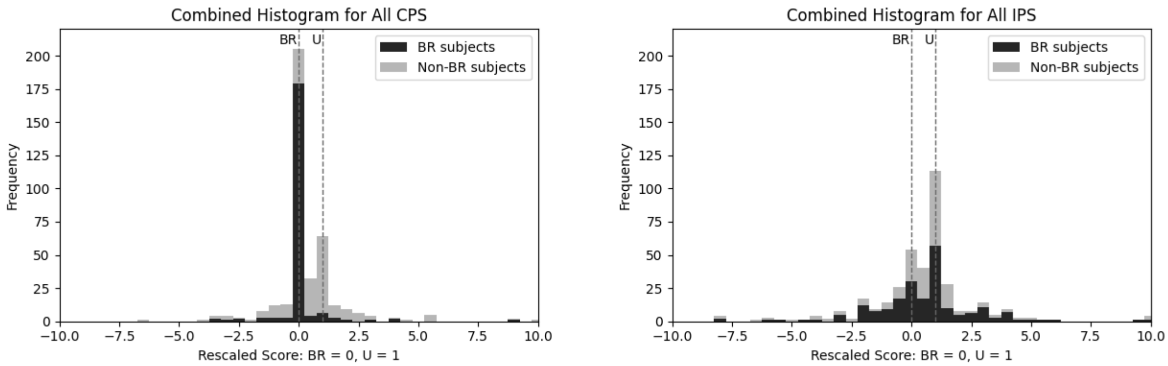

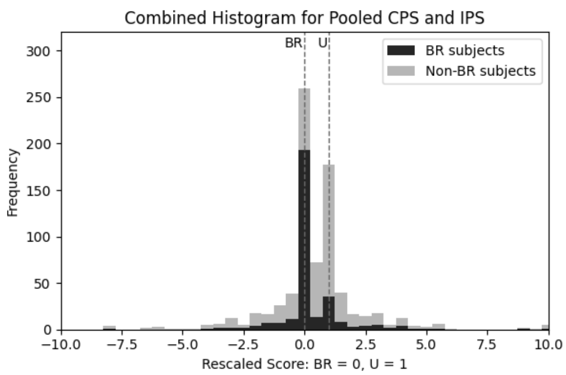

Figure Notes: This figure presents the distribution of rescaled estimates for the Conclusive, Inconclusive, and Pooled Principal Structures. Each histogram displays a stacked distribution separating BR and Non-BR subjects. The x-axis is scaled such that 0 indicates the Bayes-rational estimate, and 1 corresponds to Uniform averaging of all signals.

Figure 4: Distribution of Estimates on Principal Structures

To what extent do the subjects who succeed in the CPS problems do better in the IPS problems? For these purposes, let us classify a subject as Bayes-rational (BR) if they produce at least two estimates within ±0.05 of the Bayes-rational estimate in the four CPS problems. We simply call the complementary group Non-BR.[7] Of our 99 subjects, 53 are classified as BR, and the remaining 46 are classified as Non-BR. If we look at the number of IPS problems in which subjects are close to the Bayes-rational estimate, there is no evidence that BR subjects are better. When considering subjects' mean absolute error (MAE), BR subjects do marginally better, see Table 2, but this difference is not statistically significant (the Mann-Whitney U-test gives a p-value of 0.48). Because we do not drop any outliers in our analysis, we report these results in MAE which is more resistant to outliers than mean squared error. Additionally, we use nonparametric statistical

[7]We have considered other definitions, and the conclusions are not very sensitive to changing either the required number of correct answers, or the interval width considered to be correct.

tests such as the Mann-Whitney U test and Wilcoxon signed-rank test throughout our analysis in order to accommodate limited sample sizes. Where appropriate we also conduct t-tests and paired t-tests. In all cases, results are qualitatively similar and never result in a change of statistical significance at the 95% confidence level.

Group	Subjects	$\geq 1 \pm 0.05$	$\geq 2 \pm 0.25$	avg. MAE IPS
BR	53	4	3	1.77
Non-BR	46	9	5	2.05

Table 2: IPS Performance by CPS success

Table Notes: This table shows, for both BR and Non-BR subjects, the number of subjects who estimated at least one IPS problem within 0.05 of the rational estimate, at least two IPS problems within 0.25 of the rational estimate, and the average MAE across the four IPS problems.

We summarize with the following:

Result 2: None of our subjects solve the IPS problems precisely. Their methodology seems to be based largely on intuition, with some subjects relying on a naive average. Those who succeed in the CPS problems do no better in the IPS problems, and the population is much closer to homogeneous.

4.2.1 PRE-COR Structures

The Principal Structures suffer from an inherent identification problem. In particular, when subjects perform poorly on these tasks, it could be due to some combination of correlation neglect, precision neglect, or other biases. In order to separately identify biases that relate to correlations and precisions we use smaller, simpler structures, aimed at isolating the influence of these biases on the updating process. We implement this in the PRE-COR Structures, which consist of the Precision Sequence and the Correlation Sequence.

Referring back to Figure 3, the Precision Sequence (PRE) consists of six structures that each utilize a single computer report. In each case the Bayes-rational estimate is simply a scalar multiple of the report. The profile of rational coefficients across the six structures is $(1/4, 1/3, 1/2, 2/3, 3/4, 1)$. Note that, with the exception of PRE6, these structures are inconclusive. These structures are designed to isolate the influence of connectivity, as the signals vary in precision but are not correlated. We call a subject precision-rational (P-R) if at least 3 of their 6 estimates are within the same 0.05 radius of the Bayes-rational estimate.

The Correlation Sequence (COR) consists of eight structures, in which the subject observes a single computer report and estimates the report produced by a second computer. The Bayes-rational estimate is a scalar multiple of the observed report, and the profile of Bayesian coefficients across the sequence is $(0, 1/4, 1/3, 1/2, 1/2, 2/3, 3/4, 1)$. These structures are also inconclusive. These problems are designed to isolate the influence of correlations across signals, and we call a subject correlation-rational (C-R) if at least 4 of their 8 estimates are within 0.05 of the Bayes-rational estimate. Table 3 depicts the joint distribution of success in these PRE-COR Structures.

	C-R	C-NR	Total
P-R	43.2%	6.8%	50.0%
P-NR	15.9%	34.1%	50.0%
Total	59.1%	40.9%	100.0%

Table 3: Joint Distribution of PRE-COR Structure Rationality

It seems natural to conjecture that correlations are more difficult to understand, as they explicitly involve the interaction of multiple signals. However, we find some evidence in the opposite direction. Of the 44 subjects from the PRE-COR Treatment, 50% are P-R and 59.1% are C-R. Since our design is within subject, we can also compute the joint distribution of success in these PRE-COR Structures. Of the P-R subjects, 86.4% are also C-R, whereas 73.1% of the C-R subjects are also P-R. Because of this high correlation between P-R and C-R, only 3 of the 44 subjects are P-R and not C-R while only 7 are C-R and not P-R.

The PRE-COR Treatment allows us to directly compare a subject's understanding of signal precisions and correlations from the PRE-COR Structures with their responses in the Principal Structures. Table 4 summarizes Principal Structure performance by conditioning on subjects' precision and correlation rationality from the PRE-COR Structures. Here we see that subjects who succeed on the PRE-COR Structures perform significantly better on the CPS (the Mann-Whitney U test gives a p-value of 0.02). Naturally, this also means that subjects who succeed on these structures are also more likely to be classified as "BR" by performance on the CPS. However, this performance gap does not exist when considering performance on the IPS.

Group	# obs	# BR	avg. MAE CPS	avg. MAE IPS
P-R and C-R	19	13	1.25	1.66
Non-P-R and Non-C-R	15	3	3.26	1.62

Table 4: Mean Absolute Error (MAE) on Principal Structures

Table Notes: We are unable to comment on subjects who are P-R and not C-R or vice versa because too few subjects fell into these categories (7 and 3 respectively).

Result 3: Some of our subjects exhibit a good understanding of both differences in signal precision and correlation between signals. These subjects perform better than others on the CPS but not on the IPS.

4.3 Ex Ante Structures

We have thus far presented results from treatments that substantially generalize the domains in which correlation neglect has been previously studied, identified an important distinction between conclusive and inconclusive problems, and connected this with subjects' basic understanding of the correlations and precisions of signals. However, all of these designs face the inherent limitation that a subject is combining multi-dimensional information into a scalar report, so that the analyst can not identify how, exactly, a subject combines signals and a prior belief into an estimate, which is, after all, the main task at hand.

Since we are interested in *how* a subject produces their report, our final treatment asks the subject to assign a weight to each signal (and to the prior belief), before the signals are realized.[8] Using the notation of Section 2.1, we consider a problem (Γ, t) to be parameterized by the vector y of reports, with the Bayesian posterior expectation encoded by the function $\beta(y) = a^*$. In these problems, subjects use slider bars to assign weights $\gamma = (\gamma_0, \gamma_1, \ldots, \gamma_m)$ to the prior expectation of zero and the computer reports. Subjects are instructed that the expected payoff-maximizing way to use the reports has this linear form, in which reports that should be more influential receive weights of higher magnitude. Once the reports $y = (y_1, \ldots y_m)$ are realized, and setting $y_0 = 0$, the subject's response is computed as $a = \gamma \cdot (0, y) = \sum_{k=0}^{m} \gamma_k y_k$.

We have previously identified a set of Bayes-rational subjects who are able to compute $\beta(y^0)$ for any given y^0 in a CPS. Does this ability imply that those subjects understand the qualitative features of the function $\beta(\cdot)$? The Ex Ante

[8] This approach is analogous to the "strategy method" in game theoretic experiments, where a subject is asked to report a plan for different paths of play, even though only one will be realized.

Structures allows us to study this question by directly eliciting the weights that dictate the comparative statics (partial derivatives) of $\beta(\cdot)$.

One way to summarize the qualitative findings of this exercise is to focus on the sign (positive, negative, or zero) of the weights subjects assign, as this represents arguably the most basic understanding of the comparative statics of the Bayesian posterior estimate. Recall that the optimal weights are summarized in Figure 2.

		Prior			Positive			Zero			Negative		
	Subjects	+	0	−	+	0	−	+	0	−	+	0	−
CPS	BR	25.0%	**65.0%**	10.0%	**83.3%**	8.9%	7.8%	73.3%	**20.0%**	6.7%	61.1%	14.4%	**24.4%**
	Non-BR	52.0%	**34.0%**	14.0%	**78.2%**	5.3%	16.4%	80.0%	**8.0%**	12.0%	64.0%	9.3%	**26.7%**
	Pooled	37.3%	**50.9%**	11.8%	**81.0%**	7.3%	11.7%	76.4%	**14.5%**	9.1%	62.4%	12.1%	**25.5%**
IPS	BR	**44.2%**	44.2%	11.7%	**84.4%**	8.5%	7.0%				63.3%	6.7%	**30.0%**
	Non-BR	**42.0%**	37.0%	21.0%	**75.6%**	3.6%	20.9%				76.0%	12.0%	**12.0%**
	Pooled	**43.2%**	40.9%	15.9%	**80.4%**	6.3%	13.3%				69.1%	9.1%	**21.8%**
Pooled	BR				**83.9%**	8.7%	7.4%	73.3%	**20.0%**	6.7%	61.7%	12.5%	**25.8%**
	Non-BR				**76.9%**	4.4%	18.7%	80.0%	**8.0%**	12.0%	67.0%	10.0%	**23.0%**
	Pooled				**80.7%**	6.8%	12.5%	76.4%	**14.5%**	9.1%	64.1%	11.4%	**24.5%**

Table 5: Distribution of weight signs

Table Notes: Reports are split into three categories according to the sign of their coefficient in the Bayes-rational formula. Subjects are classified as BR or Non-BR by CPS performance. Each entry gives the distribution of signs from the elicited weights for the corresponding subject group and report type.

Table 6 summarizes the signs of elicited weights from the *ex ante* problems, as a function of the sign of the optimal weight. The data is segregated by CPS and IPS problems and by BR and non-BR subject classification. Consider as a baseline the frequency of elicited positive weights when the optimal weight is positive, which is 80.7%, pooling everything. The high success rate comes partly from the fact that a positive weight is clearly the most natural conjecture. The success rate drops to 24.5% for negative weights, and drops further to 14.5% when the optimal weight is zero. Nonetheless it is clear that in the aggregate, subjects are correctly responding to the optimal weight. The frequency of positive (negative/zero) weights is highest when the optimal weight is positive (negative/zero). The main way in which the BR subjects substantially outperform the others is that in the CPS problems, BR subjects are much more likely to correctly assign zero weight when a piece of information is not relevant (the Mann-Whitney U test gives a p-value of 0.003). This applies both to the prior beliefs and to several signals in the CPS problems.[9] .

[9] Also, BR subjects correctly assign a negative weight in the IPS problems more frequently than the others, although their performance is nearly identical looking at negative weights in the CPS problems, and there is only one signal in the IPS problems that has an optimal negative weight.

Despite clear evidence that subjects are responding to the structure of the problems, the subjects fail to match the sign more than half the time when the optimal weight is negative or zero. In this sense, even of the subjects who can compute $\beta(y^0)$ for any given y^0, when asked in what direction their answer should change if a given signal had taken a greater value, most would answer incorrectly when the correct answer is either "lower" or "unchanged." This suggests a limitation of subjects' abilities to reason about related problems in the real world, where, for example, one may want to consider how their investment strategy might change if a given forecast were to become more optimistic.

An alternative way to summarize the data is to look for "Bayes-monotone" weights, i.e., profiles of weights that have the same rank order as the Bayes-rational weights.[10] Figure 5 shows the CDFs of the number of monotone profiles (out of the eight principal structures) for BR and non-BR subjects. The distribution of the number of Bayes-monotone profiles from BR subjects stochastically dominates that of non-BR subjects, suggesting that, in aggregate, the BR subjects are somewhat better approximating the properties of Bayesian updating.

[10]More specifically, a profile $(\gamma_0, \gamma_1, ..., \gamma_m)$ of weights is called monotone if (i) for reports j, k whose optimal weights are equal, $|\gamma_j - \gamma_k| \leq 0.05$, and (ii) for reports j, k where the j's optimal weight is strictly greater than the k's optimal weight, $\gamma_j > \gamma_k$.

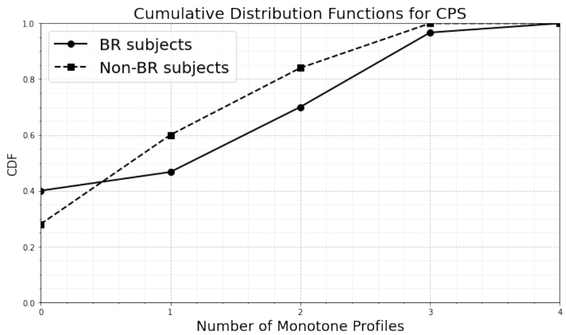

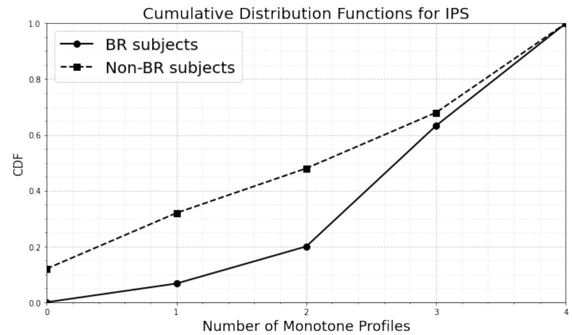

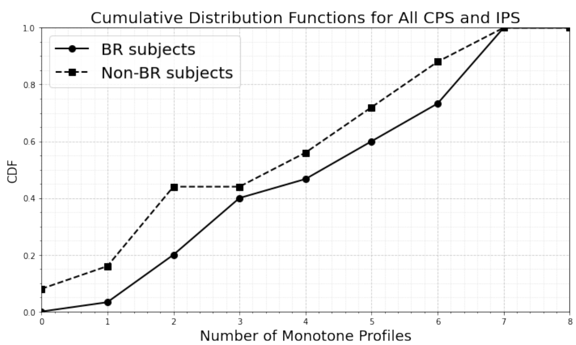

Figure Notes: This figure presents the cumulative distribution functions of the number of monotone profiles for the Conclusive, Inconclusive, and Pooled Ex Ante Principal Structures. Each solid line represents the CDF of BR subjects' monotone profiles. Each dashed line represents the CDF of Non-BR subjects' monotone profiles.

Figure 5: CDFs of the number of monotone profiles.

The implied estimates from the Ex Ante Structures are similar to the estimates from the IPS baseline; see Figure 6. In this *ex ante* form, the CPS problems do not appear to be substantially easier than the IPS problems, and subjects do not precisely provide the Bayesian estimate in either case. BR and non-BR subjects perform similarly, and the naive uniform averaging of signals is more prominent in the IPS problems. This suggests that BR subjects who correctly solve the *ex post* CPS problems are unable to extend the algorithm abstractly to essentially replace realized signal values with variables in the *ex ante* CPS problems.

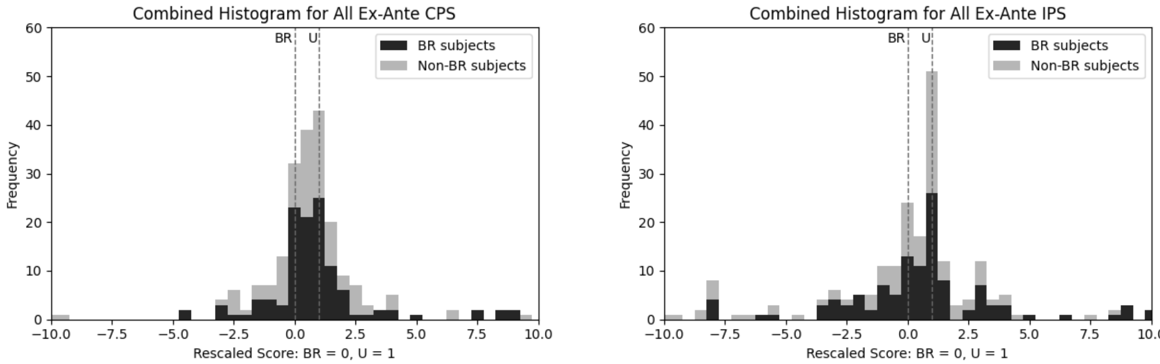

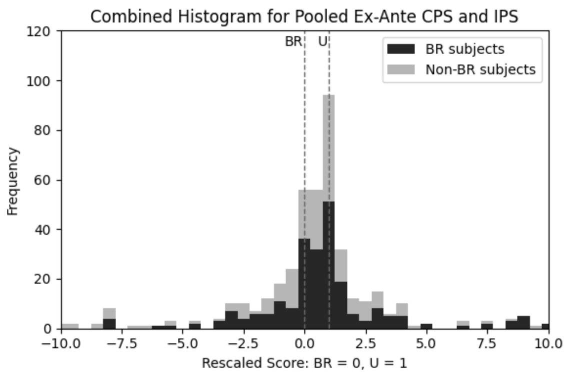

Figure Notes: This figure presents the distribution of rescaled implied estimates for the Conclusive, Inconclusive, and Pooled Ex Ante Principal Structures. Each histogram displays a stacked distribution separating BR and Non-BR subjects. The x-axis is scaled such that 0 indicates the Bayes-rational estimate, and 1 corresponds to Uniform averaging of all signals.

Figure 6: Distribution of Implied Estimates on Ex Ante Principal Structures

We summarize in the following:

Result 4: Most subjects, including those who correctly solve the baseline CPS problems, fail to consistently identify the directional effects that signals should have on their estimates. Subjects who correctly solve the *ex post* CPS problems, in aggregate, demonstrate a slightly better understanding of the properties of Bayesian updating in the *ex ante* problems.

The Ex Ante Treatment included a final phase in which subjects were presented with the same eight Principal Structures for a third time. This occurred before any feedback had been provided, and subjects were told that the random realizations were identical to those in the baseline *ex post* phase. In this third phase, subjects were reminded on their screen of (i) the realized signals, (ii) the estimate they provided in the *ex post* Phase 1, and (iii) the estimate generated by the formula they provided in the *ex ante* Phase 2.

Overall, only 10.3% of all estimates provided were outside of the interval

	Ex Post			Ex Ante			Between		
	BR	Non-BR	Pooled	BR	Non-BR	Pooled	BR	Non-BR	Pooled
CPS average	69.2%	30.0%	51.4%	7.5%	13.0%	10.0%	18.3%	40.0%	28.2%
IPS average	47.5%	29.0%	39.1%	25.0%	20.0%	22.7%	24.2%	32.0%	27.7%
Pooled average	58.4%	29.5%	45.3%	16.3%	16.5%	16.4%	21.3%	36.0%	28.0%

Table 6: Repeated session behavior in BR group and Non-BR group

Table Notes: Ex Post: subjects who report an estimate that is in the range of their estimate in the Ex Post Structures ±0.05. Ex Ante: subjects who report an estimate that is in the range of their implied estimate in the Ex Ante Structures ±0.05. Between: subjects who report an estimate between their *ex ante* and *ex post* answers but not within 0.05 of either.

bounded by the subject's Phase 1 and Phase 2 estimates (±0.05). We clearly see that the BR subjects are more likely than the Non-BR subjects to replicate their Phase 1 response (the Mann-Whitney U test gives a p-value of 0.002), and this effect comes mostly from the CPS problems. BR subjects are also significantly more likely to replicate their responses to the *ex post* CPS problems than to the *ex post* IPS problems (the Wilcoxon signed-rank test gives a p-value of 0.003). This reflects greater confidence in their Phase 1 answers to the CPS structures. These BR subjects have not only "solved" the *ex post* CPS but also know that they have solved the *ex post* CPS. In contrast, the Non-BR subjects are not significantly more likely to replicate their Phase 1 response for the CPS problems than they are for the IPS problems (the Wilcoxon signed-rank test gives a p-value of 0.79). This reflects an equal lack of confidence in their *ex post* estimates for both sets of problems. We summarize in the following:

Result 5: Subjects who correctly solve the *ex post* CPS problems are more confident in their *ex post* CPS estimates than their *ex post* IPS estimates. Subjects who do not correctly solve the *ex post* CPS problems are not more confident in their *ex post* estimates for either set of structures.

5 Conclusion

Our goal has been to analyze the way human subjects process interdependent sources of information, with a particular focus on the external validity of laboratory results. Our data is consistent with previous literature within the small portion of our experimental design that overlaps with existing work. Yet our main contributions suggest a substantially different interpretation of how well people understand these problems. Those who succeed in the baseline CPS problems seem to do so by applying a simple deductive algorithm. These subjects do not precisely solve the IPS problems. In fact, they fair no better in these problems than other subjects, and most of them assign wrong signs to the weights they give individual signals. We therefore reach more conservative conclusions about how well people process information in terms of what we think is most likely to be important outside the lab.

We hope that this work inspires future research, as there remain important issues that deserve to be explored further. We offer several thoughts on possible directions. First, if signals incorporate a noisy component, then all problems are inconclusive and the nature of updating can be qualitatively different. Redundant signals will generally retain positive marginal value and the Bayesian coefficients that govern updating can change signs. Second, we have used explicit relationships between signals and made them as transparent as possible. If subjects have imperfect information about the network structure, how will they learn about it, and what inferences will they make from realized signals over time? While we do not know of work in this direction in closely related environments, Fréchette *et al.* (2024) provide an interesting study in a different context about how subjects extract statistical information from their observations to construct a model of the data-generating process. Third, we have kept the network structure as simple as possible. When interdependencies are more complicated, involving multiple layers that relate to each other in more complex ways, how will subjects respond?

References

Agranov, Marina, and Reshidi, Pëllumb. 2024. Disentangling suboptimal updating: task difficulty, structure, and sequencing. *Working Paper.*

Arrieta, Gonzalo, and Nielsen, Kirby. 2024. *Procedural decision-making in the face of complexity.* Tech. rept. Working Paper.

Chandrasekhar, Arun G, Larreguy, Horacio, and Xandri, Juan Pablo. 2020. Testing models of social learning on networks: Evidence from two experiments. *Econometrica*, **88**(1), 1–32.

Charness, Gary, and Levin, Dan. 2009. The origin of the winner's curse: a laboratory study. *American Economic Journal: Microeconomics*, **1**(1), 207–36.

Charness, Gary, Karni, Edi, and Levin, Dan. 2010. On the conjunction fallacy in probability judgment: New experimental evidence regarding Linda. *Games and Economic Behavior*, **68**(2), 551–556.

Choi, Syngjoo, Gale, Douglas, and Kariv, Shachar. 2012. Social learning in networks: a quantal response equilibrium analysis of experimental data. *Review of Economic Design*, **16**(2), 135–157.

DeMarzo, Peter M, Vayanos, Dimitri, and Zwiebel, Jeffrey. 2003. Persuasion bias, social influence, and unidimensional opinions. *The Quarterly journal of economics*, **118**(3), 909–968.

Enke, Benjamin. 2024. *The cognitive turn in behavioral economics.* Tech. rept. Working Paper.

Enke, Benjamin, and Zimmermann, Florian. 2019. Correlation neglect in belief formation. *The Review of Economic Studies*, **86**(1), 313–332.

Esponda, Ignacio, and Vespa, Emanuel. 2014. Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, **6**(4), 180–202.

Eyster, Erik, and Rabin, Matthew. 2010. Naive herding in rich-information settings. *American economic journal: microeconomics*, **2**(4), 221–43.

Eyster, Erik, and Rabin, Matthew. 2014. Extensive imitation is irrational and harmful. *The Quarterly Journal of Economics*, **129**(4), 1861–1898.

Eyster, Erik, and Weizsacker, Georg. 2016. Correlation neglect in portfolio choice: Lab evidence. *Available at SSRN 2914526*.

Eyster, Erik, Rabin, Matthew, and Weizsacker, Georg. 2015. An experiment on social mislearning. *Available at SSRN 2704746*.

Fréchette, Guillaume R, Vespa, Emanuel, and Yuksel, Sevgi. 2024. Extracting Models From Data Sets: An Experiment. *Available at SSRN 5026751*.

Garfagnini, Umberto, and Walker-Jones, David. 2024. Naive Frequentists and Redundant Information. *Available at SSRN 4901365*.

Grimm, Veronika, and Mengel, Friederike. 2020. Experiments on belief formation in networks. *Journal of the European Economic Association*, **18**(1), 49–82.

Hanna, Rema, Mullainathan, Sendhil, and Schwartzstein, Joshua. 2012. *Learning through noticing: theory and experimental evidence in farming*. Tech. rept. National Bureau of Economic Research.

Hossain, Tanjim, and Okui, Ryo. 2024. Belief formation under signal correlation. *Games and Economic Behavior*, **146**, 160–183.

Moore, Don A, and Healy, Paul J. 2008. The trouble with overconfidence. *Psychological review*, **115**(2), 502.

Ortoleva, Pietro, and Snowberg, Erik. 2015. Overconfidence in political behavior. *American Economic Review*, **105**(2), 504–35.

6 Appendix

6.1 Proof of Proposition 1

We first show the proof for the more general case of $\boldsymbol{x} \sim N(\boldsymbol{\mu_x}, \boldsymbol{\Sigma_x})$. We then apply these results to the special case of $\boldsymbol{x} \sim N(\boldsymbol{0}, \boldsymbol{\sigma^2 I_n})$. We assume throughout that $\boldsymbol{\Gamma}$ has full row rank.

The subject's problem (P) is

$$\arg\max_{a \in \mathbb{R}} \mathbb{E}[\pi(a, \boldsymbol{t'\hat{x}})] = \int \pi(a, \boldsymbol{t'\hat{x}}) L(\boldsymbol{\hat{x}}|\boldsymbol{\hat{y}}) d\boldsymbol{\hat{x}} \tag{P}$$

where $\pi(a, \boldsymbol{t'\hat{x}}) = -(a - \boldsymbol{t'\hat{x}})^2$, and $L(\boldsymbol{\hat{x}}|\boldsymbol{\hat{y}})$ is the conditional likelihood of $\boldsymbol{\hat{x}}$ given $\boldsymbol{\hat{y}}$ under the prior information.

Proposition 2 (Generalized Proposition 1)**.** *We relate the subject's estimation problem to the following distance-minimization problem:*

$$\min_{\boldsymbol{\hat{x}} \in \mathbb{R}^n} (\boldsymbol{\hat{x}} - \boldsymbol{\mu_x})' \boldsymbol{\Sigma_x^{-1}} (\boldsymbol{\hat{x}} - \boldsymbol{\mu_x}) \quad s.t.\ \boldsymbol{\Gamma\hat{x}} = \boldsymbol{\hat{y}}. \tag{P2}$$

If sources are independent, then the solution to (P2) is given by

$$\boldsymbol{\hat{x}^*} = \boldsymbol{\Sigma_x \Gamma'} (\boldsymbol{\Gamma \Sigma_x \Gamma'})^{-1} (\boldsymbol{\hat{y}} - \boldsymbol{\Gamma \mu_x}) + \boldsymbol{\mu_x} \tag{5}$$

and the solution to (P), the optimal a^*, is given by $a^* = \boldsymbol{t'\hat{x}^*}$.

Proof. First, we transform (P) into a constrained likelihood maximization problem over the vector of sources, $\boldsymbol{\hat{x}}$. Since $L(\boldsymbol{x}|\boldsymbol{\mu_x}, \boldsymbol{\Sigma_z})$ is the density function of a multivariate normal random vector, it is symmetric and we have

$$\boldsymbol{t'\hat{x}^*} = a^*$$

where

$$\boldsymbol{\hat{x}^*} = \arg\max_{\boldsymbol{\hat{x}} \in \mathbb{R}^n} L(\boldsymbol{\hat{x}}|\boldsymbol{\mu_x}, \boldsymbol{\Sigma_x}) \quad s.t.\ \boldsymbol{\Gamma\hat{x}} = \boldsymbol{\hat{y}} \tag{P3}$$

We next show that (P3) is equivalent to (P2) under the assumption that sources are independent. Since (P3) has the same constraint as (P2), it suffices to show that (P3) and (P2) share the same objective function. Given that sources are independent, $\boldsymbol{\Sigma_x}$ is a diagonal matrix and $L(\hat{\boldsymbol{x}}|\boldsymbol{\mu_x}, \boldsymbol{\Sigma_x})$ can be written as

$$L(\hat{\boldsymbol{x}}|\boldsymbol{\mu_x}, \boldsymbol{\Sigma_x}) = \prod_{i=1}^{n} l(\hat{x}_i|\mu_i, {\sigma_i}^2) = (\prod_{i=1}^{n} \frac{1}{\sqrt{2\pi{\sigma_i}^2}}) exp(-\sum_{i=1}^{n} \frac{(\hat{x}_i - \mu_i)^2}{2\sigma_i^2})$$

where μ_i is the i-th entry of vector $\boldsymbol{\mu_x}$ and σ_i^2 is the ii-th entry of matrix $\boldsymbol{\Sigma_x}$. After eliminating the constants, it is straightforward to show that maximizing $L(\hat{\boldsymbol{x}}|\boldsymbol{\mu_x}, \boldsymbol{\Sigma_x})$ is equivalent to (P2).

$$\min_{\hat{\boldsymbol{x}} \in \mathbb{R}^n} (\hat{\boldsymbol{x}} - \boldsymbol{\mu_x})' \boldsymbol{\Sigma_x^{-1}} (\hat{\boldsymbol{x}} - \boldsymbol{\mu_x}) \tag{P2}$$

$$s.t. \ \boldsymbol{\Gamma} \hat{\boldsymbol{x}} = \hat{\boldsymbol{y}}.$$

Next, we derive the solutions to (P2) and (P) as claimed in Proposition 1. The First Order Conditions (FOCs) yield

$$2(\hat{\boldsymbol{x}} - \boldsymbol{\mu_x})' {\boldsymbol{\Sigma_x}}^{-1} + \boldsymbol{\lambda'\Gamma} = \boldsymbol{0}$$

$$\boldsymbol{\Gamma} \hat{\boldsymbol{x}} = \hat{\boldsymbol{y}}$$

From the first FOC we have

$$\hat{\boldsymbol{x}}^* = -\frac{1}{2} \boldsymbol{\Sigma_x \Gamma' \lambda} + \boldsymbol{\mu_x}$$

$\boldsymbol{\Sigma_x}$ is a diagonal matrix and $\boldsymbol{\Gamma}$ has full row rank by assumption. Therefore, $\boldsymbol{\Gamma\Sigma_x\Gamma'}$ is invertible and we have

$$\boldsymbol{\lambda} = -2(\boldsymbol{\Gamma\Sigma_x\Gamma'})^{-1}(\hat{\boldsymbol{y}} - \boldsymbol{\Gamma\mu_x})$$

Substituting yields

$$\hat{\boldsymbol{x}}^* = \boldsymbol{\Sigma_x\Gamma'}(\boldsymbol{\Gamma\Sigma_x\Gamma'})^{-1}(\hat{\boldsymbol{y}} - \boldsymbol{\Gamma\mu_x}) + \boldsymbol{\mu_x} \tag{5}$$

and the solution to (P), the optimal a^*, is given by $a^* = \boldsymbol{t'}\hat{\boldsymbol{x}}^*$ as claimed.

Finally, we show the result for the special case of $\boldsymbol{\mu_x} = \boldsymbol{0}$ and $\boldsymbol{\Sigma_x} = \sigma^2 \boldsymbol{I_n}$. In this case, (P2) reduces to (P1) and (5) reduces to (1). Furthermore, our assumption that $\boldsymbol{\Gamma}$ is full row rank is sufficient to ensure that $\boldsymbol{\Gamma\Gamma'}$ is invertible.

Therefore, we relate (P) to (P1) in this case as claimed.

$$\min_{\hat{\boldsymbol{x}} \in \mathbb{R}^n} \hat{\boldsymbol{x}}' \hat{\boldsymbol{x}} \tag{P1}$$

$$s.t.\ \boldsymbol{\Gamma} \hat{\boldsymbol{x}} = \hat{\boldsymbol{y}}$$

and the solutions are

$$\hat{\boldsymbol{x}}^{\boldsymbol{*}} = \boldsymbol{\Gamma}'(\boldsymbol{\Gamma}\boldsymbol{\Gamma}')^{-1}\hat{\boldsymbol{y}} \tag{1}$$

$$a^* = \boldsymbol{t}'\hat{\boldsymbol{x}}^{\boldsymbol{*}} \tag{2}$$

□

6.2 Formal definition of behavioral patterns

In this section we formally define four behavioral patterns.

Bayes Rational (BR): A Bayes rational agent reports the rational solution to the problem using $\hat{\boldsymbol{x}}^{\boldsymbol{*}} = \boldsymbol{\Gamma}'(\boldsymbol{\Gamma}\boldsymbol{\Gamma}')^{-1}\hat{\boldsymbol{y}}$. Such an agent is rational about connectivity, $\forall i \in \{1, ..., m\},\ \kappa_i^k = \kappa_i$; and also overlap, $\forall i, i' \in \{1, ..., m\},\ i \neq i',\ \omega_{ii'}^k = \omega_{ii'}$.

Precision Neglect (PN): An agent exhibiting precision neglect ignores connectivity differences, and is otherwise rational. Let $\kappa_{max} = \max_{i \in 1,...,m} \kappa_i$ be the highest connectivity in a given network. A precision neglector uses $\kappa_i^k = \kappa_{max}$, $\forall i \in \{1, ..., m\}$ for connectivity. A precision neglector is fully rational about overlap. $\forall i, i' \in \{1, ..., m\},\ i \neq i',\ \omega_{ii'}^k = \omega_{ii'}$.

Correlation Neglect (CN): An agent exhibiting correlation neglect sets the overlap between any two reports as zero, but is rational about connectivity. For connectivity, $\forall i \in \{1, ..., m\},\ \kappa_i^k = \kappa_i$. For overlap, $\forall i, i' \in \{1, ..., m\},\ i \neq i',\ \omega_{ii'}^k = 0$.

Uniform (U): A uniform agent ignores connectivity differences and sets the overlap between any two reports to zero. For connectivity, $\forall i \in \{1, ..., m\},\ \kappa_i^k = \kappa_{max}$. For overlap, $\forall i, i' \in \{1, ..., m\},\ i \neq i',\ \omega_{ii'}^k = 0$.