THE WISDOM OF THE CROWD AND HIGHER-ORDER BELIEFS

YI-CHUN CHEN[A], MANUEL MUELLER-FRANK[B], AND MALLESH M. PAI[C]

ABSTRACT: We propose a new simple procedure called *Population-Mean-Based Aggregation* (PMBA) that enables a principal to "aggregate" information about an unknown state of the world from agents without understanding the information structure among them. PMBA only requires agents to communicate their beliefs about the state, and some agents to communicate their expectations of the population average belief. In a large population, for any finite number of possible states, and under weak assumptions on the information structure, including individual agents' beliefs being possibly misspecified, we show that PMBA always infers the true state. We provide evidence that our procedure performs well on the experimental data of Prelec, Seung, and McCoy (2017).

KEYWORDS: aggregating beliefs, higher-order beliefs, wisdom of the crowd.
JEL CLASSIFICATION: D82, D83.

[A]DEPARTMENT OF ECONOMICS AND RISK MANAGEMENT INSTITUTE, NATIONAL UNIVERSITY OF SINGAPORE, ECSYCC@NUS.EDU.SG
[B]IESE BUSINESS SCHOOL, UNIVERSITY OF NAVARRA, MMUELLERFRANK@IESE.EDU
[C]DEPARTMENT OF ECONOMICS, RICE UNIVERSITY, MALLESH.PAI@RICE.EDU
JANUARY 27, 2024
We are especially grateful to Drazen Prelec, H. Sebastian Seung and John McCoy for access to their experimental data from Prelec, Seung, and McCoy (2017). We thank Itai Arieli, Rahul Deb, Ben Golub, Tai-Wei Hu, Yanwei Jia, Scott Kominers, Ran Smorodinsky, Tom Wilkening; and seminar/conference audiences at a16z Crypto Research, Caltech, Games 2020, IIM Ahmedabad, Bonn, IDEAL Institute at Northwestern, for helpful comments and discussions. We also thank Samuel Ho, Mingzi Niu, Hang Shao, and Xinhan Zhang for excellent research assistance. Chen gratefully acknowledges the financial support from the Singapore Ministry of Education Academic Research Fund Tier 1. Mueller-Frank gratefully acknowledges the financial support of the Spanish Ministry of Science, Innovation and Universities (Ref. ECO2015-63711-P: MINECO/FEDER, UE) and the Department of Economy and Knowledge of the Generalitat de Catalunya (Ref. 2017 SGR 1244). Pai gratefully acknolwedges support from the NSF (CCF-1763349).

1. INTRODUCTION

A long-held belief, starting from at least Aristotle, is that of the "wisdom of the crowd." This refers to the idea that even though some event may be uncertain, and the information each individual has may be noisy, the "aggregate" of individual beliefs is accurate. Such aggregation obviously has large social and private value, especially when it concerns important social or economic events. As a consequence, there has been substantial effort toward both improving existing institutions (e.g., polling), and designing and implementing new methods (e.g., prediction markets).[1] At a theoretical level, then, it is important to understand the limits of such an exercise: without specific assumptions on the nature of the information agents have, how can we aggregate it, and how much can we learn from it? In short, how wise is the crowd, and how can we tap into its wisdom?

To fix ideas, imagine a population of interest and an event whose realization is uncertain. Individual agents each see a signal that is informative of the realization. Agents also understand the informational environment, i.e., what this signal implies for their own beliefs and the implied distribution of other agents' beliefs. Even assuming away the difficulties (both logistical and strategic) of eliciting the relevant information from the population, the question we address in this paper is: Does there exist a general procedure to aggregate individuals' information without knowledge of the information structure among them?

One important insight from the existing literature is that even if the state is binary, and agents' signals are i.i.d. conditional on the state, knowing the first-order beliefs of even an infinite set of agents is generally insufficient to learn the state.[2] In particular, even in these idealized conditions, most agents may be wrong, i.e., most agents may place a posterior probability of larger than 50% on the incorrect state of the world. The main takeaway from Arieli, Babichenko, and Smorodinsky (2017) is that without further strong assumptions about the signal structure, the principal cannot learn the state of the world, even with an infinite number of agents in the population. In the terminology of econometrics, there is an "identification problem."

We answer this question by means of a simple procedure we name *Population-Mean-Based Aggregation* (herein PMBA). Our procedure uses agents' beliefs about the uncertain event as well as some agents' expectation of the average of others' beliefs. Under weak conditions on the information structure generating the agents' signals, this procedure

[1]There is also a large and influential body of literature studies the possibility of such aggregation within our current institutions, e.g., voting (see e.g., Austen-Smith and Banks (1996), Feddersen and Pesendorfer (1997)), markets (e.g., Grossman and Stiglitz (1980)) and social networks (see e.g., Golub and Jackson (2010), Acemoglu, Dahleh, Lobel, and Ozdaglar (2011)).

[2]See Prelec, Seung, and McCoy (2017) and Arieli, Babichenko, and Smorodinsky (2017).

fully aggregates the information among agents in a large population. In other words, using information from agents, solely about agents' own beliefs, and their expectation of the average belief in the population (herein referred to as population average belief), but with no other knowledge of the underlying information structure, we can fully learn the underlying state. Importantly, we show that our procedure is *robust*: it allows for aggregation of beliefs even when beliefs are misspecified at both the individual and aggregate levels. That is to say, we start from a "correctly specified" model where agents have beliefs and higher-order beliefs from a common prior and show that our PMBA procedure correctly aggregates agents' information without knowledge of the underlying common prior. We then show that this procedure continues to correctly aggregate even if individual beliefs are *arbitrarily* misspecified, as long as *in aggregate*, the misspecification is not too large, in a sense that we make precise in what follows. As a result, our procedure is a viable candidate to be used in practice.

The idea of eliciting and using higher-order beliefs is not new: In our context, we note the elegant "Surprisingly Popular" (herein, SP) procedure developed by Prelec, Seung, and McCoy (2017) (henceforth PSM), who pioneered the use of second-order beliefs for the purpose of information aggregation (however their procedure aggregates information on a narrow set of information structures—we discuss this in more detail in Section 3.1). Another close comparison is to the Surprisingly Confident (SC) procedure of Wilkening, Martinie, and Howe (2021), who also leverage the linearity of the correspondence between beliefs and higher-order beliefs, and its implication for the ordering of average beliefs and average second-order beliefs in each state (see also Theorem 1.4 of Prelec, Seung, and McCoy (2017) and contemporaneous approaches by Libgober (2021) and Prelec and McCoy (2022), described in Section 6). A key differentiator is that SP/SC are both only guaranteed to work under binary states, i.e., two possible states of the world, and require stronger assumptions more generally, as we discuss below.

In line with this literature, it is important to note that we are solely using agents' beliefs and expected higher-order beliefs. Such information has been robustly elicited in laboratory and real-world settings, and therefore, our procedure is amenable to implementation. We show in Section 3 that our procedure performs well on the experimental data of Prelec, Seung, and McCoy (2017). By contrast, theoretically feasible procedures that ask agents for more detailed information such as state-contingent beliefs (for example, the agent's expectation of the average belief in the population *conditional* on the true state of the world) may not be reliable or even feasible in practice. There is a wealth of evidence that agents struggle to report richer details (such as state-contingent beliefs).[3]

[3]For example, in Li (2017), he states: "More generally, laboratory subjects find it difficult to reason state-by-state about hypothetical scenarios (Charness and Levin, 2009; Esponda and Vespa, 2014; Ngangoué and

Similarly, we may not be able to elicit the underlying prior, i.e., we take seriously the idea that the prior/ ex-ante stage is an "as-if" assumption.[4]

1.1. *Summary of Model and Results*

Our baseline model, detailed in Section 2, considers general environments with a finite number of states of the world and many agents. We suppose that agents observe signals according to an information structure that is *unknown to the principal* but understood among the agents themselves. Theorems 1 and 2 show that, under general conditions, full aggregation can be obtained using the population's first-order beliefs and, from (at least) as many agents as there are states, agents' expectations of the average belief of the population.[5] In particular, we explain how to more fruitfully use higher-order beliefs when there are multiple states (indeed, PSM themselves point out the difficulties of extending the SP procedure when there are three or more states, and we recap the issues with illustrative examples in Section 3.1).

The basic idea is intuitive and can be described verbally, especially in the case of binary states (Theorem 1, Section 2.2). For simplicity, suppose every agent sees a conditionally i.i.d. signal about the unknown state according to an information structure known to them but unknown to the principal. Now suppose the principal elicits, from each agent, their belief about the state. By the law of large numbers, the population average belief must equal the expected belief in that state.[6] If the principal knows the expected belief in each state, then they would be done: they could compare the elicited average to the expected belief in each state.[7] However, they do not know this mapping of the state to the expected beliefs.

Suppose the principal asked two agents with different beliefs for their expectation of the average belief of the population. By the law of iterated expectations, each reports the weighted average of the expected belief conditional on the state, weighted by their (different) likelihood of the two states. This presents the principal a simple system of two linear equations in two unknowns; thus the expected beliefs in each state can be uniquely recovered from the information from these two agents. We show in what follows that this basic idea is much more applicable.

Weizsäcker, 2021). This mental process, often called "contingent reasoning," has received little formal treatment in economic theory. In subsequent work, Martínez-Marquina, Niederle, and Vespa (2019) investigate axioms governing contingent reasoning in single-agent decision problems."

[4]See for example the discussion on page 926 in Gul (1998).

[5]Section 3 discusses our rationale for choosing what information we ask agents for. Specifically, the ex-ante prior is purely a modeling device. Agents are informed, i.e., the procedure interacts with agents at the interim stage. Asking agents for their priors is not possible.

[6]This is an important point, see Remarks 3 and 4 below for further discussion.

[7]For informative signals, these expected beliefs must be different in the two states.

In particular, we consider settings where agents understand their own information, but misperceive the distribution of other agents' beliefs etc.— see Theorem 3, Section 2.4. This idea of the misperception of others' information is related to the environments considered in Eyster and Rabin (2010), Bohren (2016), and Frick, Iijima, and Ishii (2020). It is arguably a more reasonable description of the real world. We consider 2 possible states of the world, and allow that individual agents may be arbitrarily misinformed about the state contingent average beliefs of others. We only assume 1) this misspecifiation about the state contingent average beliefs of others is stochastically independent of their own belief about the unknown state, and 2) *on average*, the agents' misspecification is smaller than the difference of the state contingent average beliefs. We show that in this case a variant of PMBA still successfully aggregates agents' information. Indeed, this is the basis of the procedure that we take to the data in Section 3. Further, as we discuss there, SP is not robust to this sort of misspecification.

While our present paper is theoretical, the idea of using agents' higher-order beliefs has been successfully implemented in practice, both in the lab and in the field. Section 3 discusses more generally the rationale for the information we choose to ask agents for, and possible pitfalls in asking for richer information. Section 3.1 studies the implications on the SP procedure. In particular, we demonstrate how our analysis implies that the procedure is more generally applicable than previously understood, and how our procedure is applicable when it is not.

The original SP paper (Prelec, Seung, and McCoy, 2017) showed via lab experiments that such procedures have promise not just in theory, but also in practice. However, a reader might be concerned that our procedure uses the higher-order belief information in a more exacting way than the SP procedure and related aggregation schemes: in particular the SP procedure for binary states of the world uses only the ordinal part of this information,while we attempt a matrix inversion with these beliefs. With this in mind, we provide evidence that our procedure has real-world plausibility. In Section 3 we revisit their lab-experiments.[8] We show that our procedure performs comparably with the SP procedure even when the set of states of the world is binary (i.e., where our procedure offers no advantage over the SP procedure).

A majority of this paper considers unincentivized procedures (i.e., we are not considering the incentives of agents to report truthfully, only aggregating their reports assuming that they are). Nevertheless, our procedure naturally extends to a mechanism in which truthful reporting is a Bayes-Nash equilibrium among agents, analogous to the connection between SP and Bayesian Truth Serum—see Section 4.2 for details. In the remainder

[8]We are grateful to the authors, Drazen Prelec, H Sebastian Seung and John McCoy for sharing their original data with us.

of the paper, the use of the word elicit only refers to unincentivized elicitation, and we refer to our aggregation scheme as a procedure rather than a mechanism to clarify that there are no incentives. We also show how to adapt our method when only asking for "votes" from most agents (Proposition 5, Section 4.1), e.g., the state they perceive to be most likely (this was also the main exercise considered in Prelec, Seung, and McCoy (2017)).

In Section 5, we attempt to understand the full power and limitations of aggregation with higher-order beliefs. PMBA shows that with an infinite population, access to agents' first-order and (a few agents') second-order beliefs can suffice. We consider the case of a finite population (in particular, where signals are not conditional i.i.d.). We show that if the principal can extract the full hierarchy of agents' higher-order beliefs, they can learn the "full information posterior", i.e., the posterior of an omniscient agent who shares the agents' prior belief and sees all their individual signals, despite the principal not knowing the prior and only collecting information about posteriors. Conversely, we show that if the principal only sees a finite prefix of the hierarchy (e.g., only the first k beliefs) of a population, then there remains an "identification problem," i.e., there exist two different information structures that could both result in this exact realization. In short, higher-order beliefs cannot achieve full Bayesian aggregation in the case of a finite population unless we can ask agents for their entire hierarchy, which we consider unrealistic/ impractical. Section 6 discusses the related literature in further detail and concludes.

2. Population-Mean-Based Aggregation

In this section, we introduce our core model; we describe our Population-Mean-Based Aggregation procedure (herein PMBA) and prove its validity.

2.1. *General Model and Notation*

There is an infinite population of agents. To this end, let $N = \{1, 2, \ldots\}$ be a countably infinite set of agents, with i denoting a generic agent. There are a finite set of L states of the world $\Omega = \{\omega_1, \omega_2 \ldots \omega_L\}$. The space of feasible beliefs over Ω, denoted as $\Delta(\Omega)$, is the L dimensional simplex Δ^L.

Each agent i observes a signal in a set S with the associated sigma field $\mathcal{F}$. To allow for correlation, let us denote the common prior over $\Omega \times S^N$ by P. Given that an agent observes some signal $s \in S$, their posterior belief that the state is ω is given by $P(\omega|s)$. Let us denote the posterior beliefs of agent i by $\mu_i(s) \in \Delta^L$ and $\tilde{\mu}_i$ as the associated random variable. In what follows, the actual signal seen by the agent will often be irrelevant since we are concerned with their beliefs. We will, therefore, suppress the dependence on the signal in our notation. We define

$$\overline{\mu}_i(\omega) = \mathbb{E}[\tilde{\mu}_i|\omega].$$

In words, $\overline{\mu}_i(\omega)$ is the expected belief of agent i in state ω.

Finally, we define $\overline{\mu}(\omega) = \lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} \overline{\mu}_i(\omega)$.[9] This is the expected population average belief in state ω.

Given a vector $x \in \Delta(\Omega)$, we will let x_ω refer to the component corresponding to ω; i.e., $\mu_{i,\omega}$ is the belief of agent i that the state is ω, $\tilde{\mu}_{i,\omega}$ is the random variable of agent i's belief that the state is ω, etc.

2.2. *Binary States*

In this section, we consider the case of binary states, i.e., $L = 2$. This procedure will form the core of our general analysis.

We let $P^S_{i,\omega}$ denote the induced marginal distribution of P over signals for agent i in state ω. We also abuse notation so that μ_i represents agent i's belief that the state is ω_1: since there are only two states, their belief that the state is ω_2 is $1 - \mu_i$.

ASSUMPTION 1 (No Aggregate Uncertainty)**.** *We make the following assumptions on P:*

(1) *Imperfectly informed agents: for each agent i, we have that $P^S_{i,\omega}$ and $P^S_{i,\omega'}$ are mutually absolutely continuous with respect to each other.*
(2) *Limited correlation: For every $\varepsilon > 0$, there exists a finite $n(\varepsilon)$ such that conditional on the state, for any agent i, all but at most $n(\varepsilon)$ agents' signals are ε-independent of agent i's signal. Formally:*

$$\forall\varepsilon > 0, \exists n(\varepsilon) \in \mathbb{N} \text{ s.t. } \forall i, \exists N_i \subseteq N, |N_i| \le n(\varepsilon),$$
$$\forall j \in N\backslash N_i, \forall E, E' \in \mathcal{F}, \forall\omega : |P(\tilde{s}_{i,\omega} \in E, \tilde{s}_{j,\omega} \in E'|\omega) - P^S_{i,\omega}(E)P^S_{j,\omega}(E')| \le \varepsilon.$$

This assumption may seem a little dense so some discussion is useful. The first condition is simply to make the problem interesting/non-trivial: after all, if an agent is perfectly informed then aggregation is trivial. Next, we need an assumption regarding the correlations between agents' signals. Intuitively, assuming conditionally i.i.d. signals is likely overly strong and does not capture the idea that agents' information may be correlated. However, we cannot allow for arbitrary correlation—for example, if all agents' information is perfectly correlated, then effectively, there is only one unique opinion in the population, and no further aggregation is possible. Our notion allows for "enough independence," thereby roughly guaranteeing that for any agent, most other agents' signals are approximately independent conditional on the state. This ensures that the law of large numbers holds, so that the population average belief is a deterministic function the

[9]Note that without further assumptions, it is possible that this limit might not exist and/or might be different depending on the ordering of agents. We ignore these possibilities. Simple sufficient conditions (e.g., conditional i.i.d. signals) rule out these concerns with an infinite population, and they are moot for any arbitrarily large finite population.

of state, i.e., in the notation of our model, in any state ω we have that

$$\text{as } n \to \infty, \ \ \frac{1}{n}\sum_{i=1}^{n}\mu_i \to_P \overline{\mu}(\omega).$$

We emphasize that while this assumption relaxes the conditional i.i.d. assumption that is regularly made, it is a non-trivial assumption. Whether it is sensible or not depends on the context. For instance, in a political context, if all agents are drawing their information from the same few noisy news sources, then there will be correlations in the beliefs and our assumption would not be well-motivated for such a setting. Conversely, if we are considering subjective but idiosyncratic product reviews, then the assumption is far more natural.

As we will see in what follows, our PMBA procedure requires that the population expected beliefs differ in the two states, i.e., $\overline{\mu}(\omega_1) \neq \overline{\mu}(\omega_2)$. This is a very weak condition, but one that one may nevertheless worry about. In principle, different distributions of beliefs may nevertheless have the same expected belief. The following assumption delivers this: it requires that all agents individually are at least minimally informed. Formally, this is defined with respect to the total variation metric. This is akin to a normalization—it simply ensures that we do not have infinitely many arbitrarily uninformed agents.

To this end, we need some more notation: let $G_{i,\omega}$ denote the induced distribution of private beliefs of agent i in state ω. We will abuse notation and also let it denote the cumulative distribution function of the belief of agent i in state ω.

ASSUMPTION 2 (Minimal Information)**.** *We assume that there exists $\delta > 0$, such that for any agent i, the total variation distance between G_{i,ω_1} and G_{i,ω_2} is uniformly bounded below by δ, i.e.,*

$$\exists \delta > 0 \text{ s.t. } \forall i : \sup_{E \subseteq [0,1]} |G_{i,\omega_1}(E) - G_{i,\omega_2}(E)| \geq \delta.$$

The first main result of our paper shows the power of the PMBA procedure as described in Procedure 1.

THEOREM 1. *Suppose the information structure among agents satisfies Assumptions 1 and 2. Then the PMBA procedure (Procedure 1) correctly recovers the true state of the world almost surely.*

Before we discuss a sketch of the proof and some directions to generalize the procedure, some remarks about the procedure and the result are in order.

REMARK 1. *Assumptions 1 and 2 jointly are sufficient for PMBA to be valid, but it can be shown that they are not necessary. A full characterization is easy to describe but perhaps unsatisfying—as long as the information structure is such that the law of large numbers holds, i.e., in any*

PROCEDURE 1. Population-Mean-Based Aggregation, 2 states

(1) Elicit, from each agent $i \in N$, their belief μ_i. Calculate $\hat{\mu}$ defined as:

$$\hat{\mu} = \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \mu_i$$

(2) Select two agents, A and B, such that $\mu_A \neq \mu_B$. Elicit from each of these agents their expectation of the average posterior beliefs in the population, i.e., denoting s_i as the signal seen by agent i, elicit for each $i = A, B$:

$$\alpha_i = \lim_{n\to\infty} \mathbb{E}\left[\frac{1}{n}\sum_{j=1}^{n} \tilde{\mu}_j \middle| s_i\right].$$

(3) Calculate $\overline{\mu}(\omega_1), \overline{\mu}(\omega_2)$ as:

$$\begin{pmatrix} \overline{\mu}(\omega_1) \\ \overline{\mu}(\omega_2) \end{pmatrix} = \begin{pmatrix} \mu_A & 1-\mu_A \\ \mu_B & 1-\mu_B \end{pmatrix}^{-1} \begin{pmatrix} \alpha_A \\ \alpha_B \end{pmatrix} \tag{1}$$

Note that the right-hand side is elicited from the agents directly; and since $\mu_A \neq \mu_B$, the middle matrix is invertible.

(4) Recover the state by comparing $\hat{\mu}$ calculated in step 1 with $\overline{\mu}(\cdot)$ calculated above. The state of the world is ω s.t. $\hat{\mu} = \overline{\mu}(\omega)$.

state ω, $plim_{n\to\infty}\frac{1}{n}\sum_{i=1}^{n}\tilde{\mu}_i = \bar{\mu}(\omega)$ (which is an implication of Assumption 1); and $\bar{\mu} \neq \bar{\mu}'$ (respectively, an implication of Assumption 2), it will be clear from the proofs that PMBA will successfully recover the true state of the world. Expressing this (in particular, concentration) in terms of primitives, however, appears out of reach.

REMARK 2. *It is worth re-emphasizing that the PMBA procedure does not require any knowledge of the underlying information structure among agents—a different way of interpreting our result, therefore, is that as long as these conditions are satisfied, knowledge of the information structure among agents is not needed to fully aggregate their information. That said, it is possible to construct examples where these conditions are violated but a principal who knows the information structure among agents can aggregate their information. For example, one can construct settings, e.g., with two states, such that $\bar{\mu}(\omega_1) = \bar{\mu}(\omega_2)$, but the population distribution of posteriors in these two states are different. A Bayesian principal who knows this would be able to perfectly aggregate agents' information, but PMBA would clearly fail in such a setting.*[10]

REMARK 3. *It is important, both conceptually and technically, to note that our assumptions (in particular, Assumption 1 part (2)) imply that $\bar{\mu}(\omega)$ is a degenerate random variable, i.e., a*

[10]As a simple illustration, consider the following example: there are two states, $\{\omega_1, \omega_2\}$, and two signals, $\{s_1, s_2\}$. All agents' prior is that either state is equally likely. In state ω_j, agent i receives the "correct" signal s_j with probability $\frac{1}{2} + e^{-i}$, and with the remaining probability receives the incorrect signal s_{3-j} (this clearly violates Assumption 2). PMBA will fail in this setting: the expected population average belief in both states is $\frac{1}{2}$. However, by observation, a Bayesian who knows the signal structure can learn from the agents' beliefs.

deterministic function of the state. Note that our PMBA procedure would fail otherwise because it recovers $\mathbb{E}[\bar{\mu}(\omega)|\omega]$, which may not equal the realization if $\bar{\mu}(\omega)$ is a non-degenerate random variable. One can construct settings where PMBA would fail, but a Bayesian principal who knew the information structure among agents would learn the true state of the world (for example, if $\bar{\mu}(\omega)$ is a non-degenerate r.v. but has disjoint support across the states).

Conceptually, our assumption puts additional structure on what we mean by a "state of the world." For example, consider a setting where agents' signals are exchangeable conditional on the state. In this case, DiFinetti's theorem tells us that the signals are i.i.d. conditional on state and an additional latent random variable. Our assumption then rules out the existence of additional latent random variables on which $\bar{\mu}(\omega)$ depends—the description of the state must be rich enough so that $\bar{\mu}(\cdot)$ is a deterministic function of ω.

A formal proof is given in Appendix A. Here, we want to discuss how the various assumptions play a role in the result. As we pointed out earlier, Part (1) of Assumption 1 is mainly a normalization. The key processing step in this procedure is Step 3. By the law of iterated expectations, an agent i with beliefs μ_i will report, in step (2),

$$\alpha_i = \mu_i \overline{\mu}(\omega_1) + (1 - \mu_i)\overline{\mu}(\omega_2).$$

Our procedure inverts this relationship to learn the unknown mapping from ω to $\overline{\mu}(\omega)$, using (1).

Step (4) then recovers the state by comparing the calculated $\overline{\mu}(\cdot)$ to the elicited $\hat{\mu}$. Part (2) of Assumption 1 ensures that an appropriate law of large numbers holds so that in state ω, $\hat{\mu} = \overline{\mu}(\omega)$, almost surely. Assumption 2 is used to guarantee that $\overline{\mu}(\omega_1) \neq \overline{\mu}(\omega_2)$. Therefore, the elicited $\hat{\mu}$ will almost surely match $\overline{\mu}(\omega)$ for a unique state, and the procedure thus successfully concludes. Taken together, therefore, these give us our result.

To be clear, the model and result of this section are explicitly written for a common prior model where the agents perfectly understand the signal structure among themselves, and only the principal does not. This knowledge of each other's information among agents may seem daunting or even implausible in applied contexts.

To that end, our procedure does not require agents to have an exact understanding of the distribution of signals among fellow agents. We can observe that it is sufficient for each agent to understand the aggregate implications perfectly; that is to say, the expected population beliefs in each state, $\overline{\mu}(\cdot)$ since that is all they need to know (other than their own beliefs) to calculate the information that we elicit. This is a substantially lower informational requirement. Indeed, in Section 4.1, we show how even less demanding information may be elicited (e.g., asking most agents what state they think is more likely).

However, even this may seem high—perhaps agents don't perfectly understand even the aggregate implications of the model—for example, each agent has noisy information about the true $\overline{\mu}(\cdot)$, rather than exact knowledge of it. In Section 2.4, we show how modifying our procedure allows for population aggregation in such a setting.

2.3. *$L > 2$ States*

We now consider the case of $L > 2$ states. Intuitively, we generalize the "matrix inversion" ideas of the previous subsection to propose a PMBA procedure that works for L states. However, having more than 2 states introduces some difficulties; notably, the assumptions that guarantee this process works become slightly more demanding. Assumption 1 still applies as written; observe that nothing in that assumption was specific to the case of binary states.

However, a direct generalization of Assumption 2 to more than two states is generally not sufficient to assure that the state-dependent averages are unequal across states. So, we assume the L state generalization of its implication directly.

ASSUMPTION 3. *We assume that P is such that for any two states ω and ω', $\overline{\mu}(\omega) \neq \overline{\mu}(\omega')$.*

Note that while Assumption 3 is not directly stated in terms of the primitives of our model (P), it should be clear that this is a property that can be easily satisfied.[11]

Finally, in the 2−state case, we are guaranteed the a.s. existence of two agents who have different beliefs. The appropriate generalization turns out to be the existence of L agents whose beliefs constitute a full rank $L \times L$ matrix so that the analog/generalization of (1) in step (3) can be carried out. The following assumption guarantees this:

ASSUMPTION 4. *We assume that P is such that for any agent i, the support of beliefs supp($\tilde{\mu}_i$) contains L distinct points, such that those L beliefs, viewed as L-dimensional vectors, constitute a set of full rank. Alternately and equivalently, we require that the convex-hull of the support has an interior relative to Δ^L.*

One might wonder why such a condition was not required in the case of binary states. The answer is that when $L = 2$, this condition simply reduces to requiring that the support of the agents' beliefs has at least two points in any state. Part (1) of Assumption 1, however, implies that the support of the agents' beliefs is the same across states. If this support was degenerate at a single point, we would clearly violate Assumption 2. Note that with $L > 2$, there clearly exist settings that violate Assumption 4 despite satisfying Assumptions 1 and 2; for instance, consider a setting where $L = 3$ so that the simplex is a triangle, but there are only two possible signals, so that there only two induced beliefs.

[11]For example, in the case of finite conditionally i.i.d signals, this property is generically satisfied in the belief space, in the sense of full measure.

We should note that this condition is weak in at least two ways. First, in the standard genericity sense: a generically chosen set of feasible posteriors will be full rank. Second, and perhaps more interestingly, even in settings where there are fewer signals than states (e.g., agents can run only a unique binary-outcome experiment whose outcome determines their posterior), as long as the agents can acquire enough independent signals (e.g., conduct a sufficient number of independent experiments), the set of feasible posteriors is full rank. We formally state this claim below— the proof essentially repurposes a lemma of Fu, Haghpanah, Hartline, and Kleinberg (2021).

Claim 1. *Suppose that signals are conditionally i.i.d. draws from a finite set S, and Assumption 3 is satisfied. If each agent gets $L-1$ independent draws from the distribution P^S_ω, then Assumption 4 is satisfied with the signal space S^{L-1}.*

We are now in a position to describe our main procedure, Population-Mean-Based Aggregation (PMBA) (Procedure 2 below) and state our main result for this setting.

Procedure 2. Population-Mean-Based Aggregation, L states

(1) Elicit, from each agent $i \in N$, their belief μ_i. Calculate $\hat{\mu}$ defined as:

$$\hat{\mu} = \lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} \mu_i$$

(2) Select L agents $I = \{i_1, i_2, \dots, i_L\}$ such that the $L \times L$ matrix $\boldsymbol{\mu} = (\mu_{i_1}^T, \mu_{i_2}^T, \dots, \mu_{i_L}^T)^T$ is full rank.

(3) Let s_i be the signal seen by agent i. Elicit from each of these agents their expectation of the average posterior beliefs in the population, i.e., elicit for each $i \in I$:

$$\alpha_i = \lim_{n\to\infty} \mathbb{E}\left[\frac{1}{n}\sum_{j=1}^{n} \tilde{\mu}_j \middle| s_i\right].$$

Denote $\boldsymbol{\alpha} = (\alpha_{i_1}^T, \alpha_{i_2}^T, \dots, \alpha_{i_L}^T)^T$.

(4) Solve

$$\bar{\boldsymbol{\mu}} = \boldsymbol{\mu}^{-1}\boldsymbol{\alpha} \tag{2}$$

Note that the right-hand side is elicited from agents directly and therefore $\bar{\boldsymbol{\mu}}$ can be recovered. The state ω is the row of the recovered $\bar{\boldsymbol{\mu}}$ that equals the calculated $\hat{\mu}$.

Theorem 2. *Suppose $L > 2$ and the common prior P satisfies Assumptions 1, 3, and 4. Then the PMBA procedure (Procedure 2) recovers the true unknown state of nature ω almost surely.*

2.4. *Aggregation in Misspecified Information Settings*

In practice, agents may have limited and individually incorrect knowledge about the aggregate information structure: in particular, a standard concern is that while agents may have meaningful information that shapes their *personal* beliefs, they nevertheless may be less informed about the belief of others, or even summary statistics of the beliefs of others like the average belief as elicited by our procedure.

Secondly, in practice there will be only a finite number of agents, and therefore it is unlikely that one of the probabilities recovered from the matrix-inversion will exactly match the population average beliefs.

Nevertheless similar ideas can be used to aggregate population information robustly in such settings. In the interests of brevity, we consider the case of two states; the extension to L states is analogous to the previous section.

We maintain the assumptions on the true signal generating process P as in Section 2.2, i.e., Assumptions 1 and 2. Instead of assuming common knowledge of P, we assume that the agents' information about P is misspecified—agent i believes that the state-dependent population averages are $(\alpha_i^{\omega_1}, \alpha_i^{\omega_2}) \in [0,1]^2$. Note that each agent understands how to interpret their own signal correctly, i.e., the misspecification is purely about knowledge of the population averages. We make the following assumption on the relation between agents' beliefs and the true state-dependent population means.[12]

ASSUMPTION 5. *For each agent i and state $\omega \in \Omega$ we have that*

$$\alpha_i^{\omega} = \bar{\mu}(\omega) + \zeta_i^{\omega}$$

where ζ_i^{ω} is independent of agent j's signal s_j conditional on the state, for all agents j. Further we assume that

- *Conditional on ω the ζ_i^{ω}'s satisfy limited correlation across agents.*
- *For each $\omega \in$ Omega, the ζ_i^{ω}'s are have a common mean, $E[\zeta_i^{\omega}] = E[\zeta^{\omega}]$ for all i.*
- *The distributions of α_i^1 and α_i^2 have disjoint support for all i.*

Let us briefly discuss Assumption 5. One implication of this is that each agent understands that conditional on the state their expectation of the population mean is independent of their first-order belief. The assumption that supports are disjoint is necessary,

[12] Even this weaker assumption may be violated in the real world where higher-order belief data is often noisy and boundary reports are common. That said, in practice, it may be more robust to use a linear regression (or a tobit regression to deal with boundary issues) on the lines suggested in Palley and Satopää (2020) and Peker and Wilkening (2023). We do not pursue this line in this paper as we feel it is out of scope, and further even the current mechanism shows promise on real world data as we show in our exercise in Section 3. Nevertheless, this we believe this presents a tantalizing possibility for future work.

keeping in mind that agents are fully Bayesian in their inferences, and they have misspecified knowledge. To be precise, this assumption avoids that their expectation for the population average in state ω_1 is larger than in state ω_2 which would contradict Bayes rationality. The limited correlation is necessary for a law of large numbers to hold.

Importantly note that we do not assume that $\mathbb{E}[\zeta^{\omega}] = 0$: agents' beliefs about the population averages may be biased even on average (as has indeed been observed in practice), and these biases may be different in different states.[13]

PROCEDURE 3. Misspecified Information Population-Mean-Based Aggregation

(1) Elicit, from each agent i, their belief μ_i. Calculate $\hat{\mu}$ defined as:

$$\hat{\mu} = \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \mu_i$$

(2) Elicit from all agents their interim expectation of the average posterior belief over all agents—denote this as α_i for agent i.

Partition agents into two infinite groups A and B such that $\mu_A \neq \mu_B$ (μ_A, μ_B are defined below). Calculate the average of the α_i's in each group, A, B. For example, without loss, rename agents so that $A = \{1, 2, 3, \ldots\}$. Then define

$$\alpha_A = \lim_{n\to\infty} \frac{1}{n} \sum_{i\leq n} \alpha_i,$$

$$\mu_A = \lim_{n\to\infty} \frac{1}{n} \sum_{i\leq n} \mu_i.$$

and define μ_B, α_B analogously.

(3) Calculate $\overline{\mu}(\omega)$ as

$$\begin{pmatrix} \overline{\mu}(\omega_1) \\ \overline{\mu}(\omega_2) \end{pmatrix} = \begin{pmatrix} \mu_A & 1-\mu_A \\ \mu_B & 1-\mu_B \end{pmatrix}^{-1} \begin{pmatrix} \alpha_A \\ \alpha_B \end{pmatrix} \tag{3}$$

Recover the state as:

$$\omega^0 = \arg\min_{\omega\in\Omega} |\hat{\mu} - \bar{\mu}(\omega)|. \tag{4}$$

We can now state our aggregation result in the misspecified information setting.

THEOREM 3. *Suppose the true signal distribution process satisfies Assumptions 1 and 2, and that the agents' knowledge regarding the information structure satisfies Assumption 5. Suppose further that*

$$\left|E[\zeta^1]\right| + \left|E[\zeta^2]\right| < \bar{\mu}\left(\omega_1\right) - \bar{\mu}\left(\omega_2\right). \tag{5}$$

[13] See also Remark 5 which clarifies how SC/ SP may fail in this setting.

Then, Misspecified Information Population-Mean-Based Aggregation *(Procedure 3) correctly recovers the true state of the world almost surely.*

Theorem 3 establishes that indeed the conceptual idea of second-order elicitation successfully overcomes the limited understanding of agents' regarding the private belief distributions of their peers. Compared to the PMBA aggregation under complete knowledge of the information structure, the two infinite groups A and B play the roles of the two agents in the PMBA.

Thus moving to a more realistic assumption of an incomplete and individually incorrect understanding of the information structure necessitates the elicitation of the expected population averages from large groups: we are relying on the crowd for both first- and second-order statistics. The main issue is that individual agents' understanding of the population averages is incorrect. However, by the law of large numbers we can appropriately average these individual agents' reports over two groups, A and B. The average first-order belief, and average report of population average, of agents in each group, will serve as the analogs of the two individuals in the baseline PMBA procedure (Procedure 1).

The basic intuition is that even though individual agents are misinformed about the distribution of their fellow agents' beliefs, on average these errors are not "large" (even though individual agents' errors can well be) relative to the true difference between the two population averages (recall condition (5) in Assumption 5 makes this possible). This ensures that the "closer" recovered $\bar{\mu}(\omega)$ to the realized average population distribution $\hat{\mu}$ is indeed the one corresponding to the state.

Finally, since we are using the "closest" distribution rather than an exact match, this procedure always selects a state (for the earlier procedures, with high probability under a finite population, the procedure would simply fail). It should be clear that under sufficient concentration, a finite population, high probability version of Theorem 3 holds. The exact bounds/guarantees would of course depend on assuming more details of the information structure among agents (e.g. how much correlation), so we do not pursue this avenue further in this paper. Instead, in what follows in Section 3 we report the results of taking this mechanism to the data.

3. PMBA-STYLE PROCEDURES COMPARED TO SP

We now compare our PMBA procedure to the current gold standard, i.e., the Surprisingly Popular (SP) procedure of Prelec, Seung, and McCoy (2017). First we describe SP formally, and explain formally when our procedure can potentially succeed while SP would not. Having described the theoretical properties of PMBA procedures, the natural question for a reader is while this works in theory, could it work in practice. To answer this in the affirmative, we revisit the experiment of Prelec, Seung, and McCoy (2017). We were graciously provided the raw data for their experiment by the authors to benchmark our own procedure.

3.1. *The Surprisingly Popular procedure*

The original SP procedure considers the case analogous to what we consider in Section 4.1—eliciting from each agent a "guess" or "vote" of the state they consider most likely given their beliefs. It also considers the case of directly eliciting beliefs (see Section 1.4 of the Supplementary Information Section of Prelec, Seung, and McCoy (2017)). More precisely, Section 1.4 of PSM departs from the idea of SP to identify the actual population signal distribution by eliciting the agents' posterior beliefs and their prediction of other agents' signal. Our discussion of SP here also applies to the alternative procedure proposed by Wilkening, Martinie, and Howe (2021), Surprisingly Confident (SC).[14] Like SP, SC aggregates when there are only two states; moreover, like PMBA, SC is also based on the linearity of the mapping from first-order belief to its corresponding second-order belief.

We formally define the procedure in our general environment, generalize the existing results, and highlight its limitation relative to PMBA. To start, consider the case of binary states and let the common prior P among agents satisfy Assumption 1 and Assumption 2. In our notation, the SC procedure can be described as follows:

PROCEDURE 4. Surprisingly Popular (SP), 2 states

(1) Elicit from each agent $i \in N$, their belief μ_i. Calculate $\hat{\mu}$ defined as:

$$\hat{\mu} = \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \mu_i$$

(2) Select one agent, i, and elicit their expectation of the average posterior beliefs of the population, i.e., denoting s_i as the signal seen by agent i:

$$\alpha_i = \lim_{n\to\infty} \mathbb{E}\left[\frac{1}{n}\sum_{j=1}^{n} \tilde{\mu}_j \middle| s_i\right].$$

(3) Compare the realized population average belief $\hat{\mu}$ with agent i's expectation α_i. If $\hat{\mu} > \alpha_i$, the state of the world is ω_1, otherwise it is ω_2.

This procedure is arguably even simpler than our PBMA procedure. Only one agent is required to report their expectation of the population average and no calculation is

[14]SC was named as such because PSM in their original paper referred to the subjects' posterior beliefs as their confidence. PSM's main object of analysis uses agents' guesses or votes, and their expectation of votes of others, while SC directly uses agents' posteriors and expected population-average posteriors.

required. Nevertheless, SP can recover the truth in the binary state setting as the following result shows.

PROPOSITION 1. *Consider a binary state space. Suppose that the information structure among agents satisfies Assumptions 1 and 2. Then the* SP *procedure (Procedure 4) correctly recovers the true state of the world.*

Note that this result provides a stronger theoretical guarantee of the applicability of the SP procedure: the original PSM paper considers only a finite set of possible signals. We omit a formal proof. At a high-level, this proceeds straightforwardly from Lemma 2: In a two-state setting (with two or more signals), for any agent i, the induced distribution of beliefs on ω at state ω first-order stochastically dominates the induced distribution of the belief on ω at state $\omega' \neq \omega$. Therefore the population average belief of state ω in ω must be higher than the population average belief of state ω in state ω'. Any agent's expected population average belief must be submitting a convex combination of these two. Therefore if the true state is indeed ω, it will be "surprisingly popular" relative to this expectation of the average beliefs in the population.

In summary, the SP procedure aggregates beliefs in the binary state setting. However, it is not applicable beyond two states. Indeed the original SP procedure cannot be applied to more than binary states without strong additional assumptions. The reason is that, in general, more than two states can be "surprisingly popular", leaving an indeterminacy in how to define the procedure. Even when defining the SP procedure as e.g., picking the state that surpassed by most its expectation, the procedure does not achieve aggregation. PSM impose an additional assumption (see e.g., their Theorem 3), one of "diagonal dominance," roughly:[15] the posterior of an agent who thinks state ω is most likely must also place a higher posterior on ω than any voter voting for $\omega' \neq \omega$ places on ω. We exhibit in Appendix C a counter-example showing that such an assumption is necessary.

At this point, a reader may wonder about whether other iterative or sequential procedures that work for 2 states can successfully aggregate information in settings with more than two states. The short answer is no, for similar reasons as the binary state PMBA cannot be used sequentially—recall Remark 3.

REMARK 4. *One possibility a reader may consider feasible is, in a setting with $L > 2$ states $\{\omega_1, \dots, \omega_L\}$, to run SP sequentially. That is to say first elicit beliefs corresponding to states ω_1 and its complement (states $\{\omega_2, \dots, \omega_L\}$), and run SC/SP on this "binary state" environment. However, note that in this case, the realized population average belief conditional on the second*

[15] To be clear, the setting there is one of prediction normalized votes, we are translating to the analogous assumption in our setting for convenience.

"state" ($\{\omega_2, \ldots, \omega_L\}$) is not deterministic as a function of this "state" (i.e., it depends on which of the states ω_1 thru ω_L actually realized). As a result the existing procedures do not apply.

Not to belabor the point, but this has two takeaways. Firstly, consider what is often modeled as a binary state space setting: e.g. we wish to aggregate information on the true/false question "Is Chicago the capital of Illinois". The true underlying state that determines the distribution of beliefs etc may nevertheless be multivalent (e.g. "Chicago is the capital", "Springfield is the capital", "Evanston is the capital" etc.). If that is indeed the case, existing approaches may fail even if the question of interest ("Is Chicago the capital of Illinois") takes only a binary value. PMBA can in principle be applied to such settings, with the proviso that the set of possible states needs to be correctly identified, and corresponding beliefs elicited.

Finally, one might ask whether the existing procedures are robust to misspecification, i.e., when agents' beliefs about the population average are incorrect (recall Section 2.4). When agents' information about the population average is misspecified, SP as written can straightforwardly fail (if the agent who is asked for their expected population average belief is misspecified, then the wrong state may be surprisingly popular). We follow the path taken by PSM in their implementation, and consider "average SP", i.e., ask all agents for the expected population average, average this, and then given the true population average check which state is "surprisingly popular" relative to that. We show by example that there exist settings where the conditions of Theorem 3 are satisfied, so that misspecified Information PMBA can aggregate information, but SP does not.

Remark 5. *Suppose we have that Assumption 5 is satisfied and further that $\overline{\mu}(\omega_1) = \frac{3}{4}$ and $\overline{\mu}(\omega_2) = \frac{1}{4}$, i.e. the average beliefs in the population are $\frac{3}{4}$ in state ω_1 and $\frac{1}{4}$ in state ω_2. Further suppose that $E[\xi_2] = 0$ but $E[\xi_1] = \frac{3}{16}$. In other words, on average, agents' information about the population average in state 2 is correct, but on average, agents believe that in state ω_1, the average belief in the population is $\frac{15}{16}$.*

By a simple calculation, in state ω_1, the average expected population belief is $\frac{49}{64}$. This is larger than the population average belief of $\frac{3}{4}$ in ω_1, so state ω_2 is surprisingly popular in this state, and therefore SP does not correctly aggregate information in this state.

Further, it is easy to verify that the condition of Theorem 3 holds in this case, and therefore, our misspecified Knowledge PMBA procedure correctly aggregates information.

Note that the converse is also possible, i.e. there exist settings where agents have misspecified information in which SP correctly averages state, but which violate our Assumption 5. SP correctly aggregates state whenever the average of agents' expectations of population average belief lies in the interval of population average beliefs.

3.2. *Comparing the Procedures in the Data*

To recall, the authors in Prelec, Seung, and McCoy (2017) report the results of using SP to aggregate information in three experiments, namely:

(1) *Study 1: State Capitals* Respondents got prompts of the form "X is the capital of State y". Agents had to respond with T/F, and their predicted population average.
(2) *Study 2: General Knowledge* Respondents were given general knowledge statements, e.g., "Japan has the highest life expectancy in the world." Responded with T/F, and their predicted population average.
(3) *Study 3: Dermatology* A collection of dermatologists were recruited and shown pictures of lesions. They were asked asked to evaluate whether it was benign or malignant (1–6 scale) and what they thought other dermatologists would say (1–11 scale).

Note that in each case the experimenter knows the "ground truth" (in the first two these are easily ascertained, in the latter study the pictures were chosen such that the true malignancy of the tumor was known). Therefore the experimenter can compare whether the aggregation resulting from the procedure results matches the ground truth. This can also be compared with simpler procedures such as population average or majority.

Table 1 reports the results of PMBA on the raw data from these studies, and how these results compare to SP. Formally, we used the procedure described in Section 2.4, i.e. Procedure 3, since that is best suited for "real world" use. We emphasize that our procedure, summarized in Table 1, performs roughly as well as SP.

Note that this remarkable performance of PMBA is despite the studies being in binary settings where our procedure has no theoretical advantage, but as we described earlier may be at a practical disadvantage (as we said earlier, PMBA-style procedures use the cardinal nature of the beliefs since they attempt a matrix-inversion step). Further, Prelec, Seung, and McCoy (2017) does not elicit agents' beliefs and uses their votes as we described above. However, they do ask agents to report their confidence in their assessment. We interpret this self-reported confidence as agents' first-order beliefs to perform the matrix inversion step. This achieves good results as shown in Table 1.[16]

Formally, for each experiment we conduct the following procedure, which is an amalgam of Procedures 3 and 5

(1) As we described earlier, we use the reported confidence of each agent as their belief.
(2) To determine the groups we average over (recall Procedure 3), we split the population into two groups, according to their vote on the question.

[16]Documented code has been submitted with the submission.

Study	Majority	SP	PMBA
Capitals	62%	90%	88%
Trivia	56.2%	62.5%	60.29%
Dermatology Lesions	67.5%	73.75%	65.15%

TABLE 1. Table summarizing the results of PMBA performance on Experimental data from Prelec, Seung, and McCoy (2017). Study column denotes which study is under consideration. Majority column reports the fraction of questions in that study where the majority matched the ground truth, analogously for SP and PMBA.

(3) For each group, we use the average confidence of members as the corresponding "belief", and the average of the vote shares predicted by members of the group as the corresponding predicted vote share for that group.
(4) We then solve for the predicted state-contingent expected vote shares given these according to PMBA (recall (3)).
(5) The predicted state is the one whose vote share is closer to this predicted vote share (3).

We should note two particular difficulties with our implementation: firstly, agents' confidence in this experiment is a self-report, it is not elicited in an incentivized fashion. This causes no difficulties for Prelec, Seung, and McCoy (2017) in their analysis since this is not an input to SP (recall they are using agents' "votes" as we described above), they only use it to compare SP against a confidence-weighted majority benchmark. However, for our procedure, we require (some) agents' first-order beliefs for the matrix-inversion step. Even granting this, a particular difficulty is on the third experiment: it is particularly problematic to interpret confidence as belief since elicitation was on $1-6$ and $1-11$ scales. It is an exciting prospect to design and test PMBA directly (with elicited beliefs and population average beliefs) in settings with more than two states. We defer that to future work.

4. EXTENSIONS

In this section, we show how to extend the basic PMBA to other settings. In Section 4.1, we recognize that, in practice, eliciting beliefs from agents is difficult—we show that a variant of the procedure can aggregate information if we only elicit "simple" information (e.g., report which state is thought more likely) from most agents, and elicit more fine-grained information from only a small number of agents. Both extensions are displayed for the case of 2 states of the world.

4.1. *Aggregating Guesses instead of Beliefs*

In this section we discuss how our PMBA procedure can be adjusted to environments where not beliefs but rather finite guesses are elicited, e.g., which state is more likely rather than what is the probability of state $\omega = 1$ (this is also the formulation in Prelec, Seung, and McCoy (2017)—see Section 3.1). Consider a binary state space and an information structure that satisfies Assumptions 1 and 2. The action-based PMBA procedure is defined in Procedure 5.

PROCEDURE 5. Action-based PMBA, 2 states

(1) Elicit from each agent $i \in N$, the state $\omega \in 0,1$ they consider more likely, $a_i \in 0,1$. If their belief equals 0.5, have them report $a_i = 1$.[a] Calculate $\hat{\alpha}$ defined as:

$$\hat{\alpha} = \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} a_i$$

(2) Select 2 agents, A and B, such that $a_A \neq a_B$. Elicit from each of these agents $i = A, B$ their belief μ_i and their expectation of the average report in the population, i.e., denoting s_i as the signal seen by agent i, elicit for each $i = A, B$:

$$\alpha_i = \lim_{n\to\infty} \mathbb{E}\left[\frac{1}{n}\sum_{j=1}^{n} a_j \middle| s_i\right].$$

(3) Calculate $\bar{\alpha}(\omega_1), \bar{\alpha}(\omega_2)$ as:

$$\begin{pmatrix} \bar{\alpha}(\omega_1) \\ \bar{\alpha}(\omega_2) \end{pmatrix} = \begin{pmatrix} \mu_A & 1-\mu_A \\ \mu_B & 1-\mu_B \end{pmatrix}^{-1} \begin{pmatrix} \alpha_A \\ \alpha_B \end{pmatrix} \tag{6}$$

Note that the right-hand side is elicited from agents directly; and since $\alpha_A \neq \alpha_B$, the middle matrix is invertible.

(4) Recover the state by comparing $\hat{\alpha}$ calculated in step 1 with $\bar{\alpha}(\cdot)$ calculated above. The state of the world is ω s.t. $\hat{a} = \bar{a}(\omega)$.

[a] Note that this tie-breaking cannot be ensured by an incentive-compatible elicitation mechanism.

For the action based PMBA procedure to be viable we require the existence of two agents A, B with differing first-order guesses, $a_A \neq a_B$. Assumptions 1 and 2 are not sufficient to assure this. The following assumption, however, guarantees the existence of two such agents.

ASSUMPTION 6 (No Herding)**.** *We assume that there exists $\delta > 0$, such that for any agent i and any state $\omega_j = \omega_1, \omega_2$, we have that the probability agent i assigns to state ω_j conditional on their signal s_i is larger than half with a probability of at least δ.*

We can now present our result on aggregation in this environment.

PROPOSITION 2. *Consider a binary state space. Suppose that the information structure satisfies Assumptions 1, and 6. Then the* Action Based PMBA *procedure (Procedure 5) correctly recovers the true state of the world a.s.*

The result follows along the same lines as Theorem 1.

4.2. *Incentivizing Truthful Reporting*

In our analysis so far we abstracted away from incentive issues, assuming truth-telling. We now discuss how truthful elicitation and aggregation can simultaneously be achieved via a slightly adjusted mechanism that includes payments which ensure truthful reporting is a (strict) Bayes Nash Equilibrium.

Recall that when there is a verifiable state that is or will be known to the principal, the principal can truthfully elicit agents' beliefs about the state by paying them using a proper scoring rule. A proper scoring rule pays agents as a function of the reported beliefs and the true state so that for any agent, truthful reporting of beliefs is a dominant strategy. For example, classic scoring rules such as the Brier score (Brier, 1950) or the logarithmic scoring rule make truthful reporting a dominant strategy for expected utility maximizing agents.[17]

When the state is not verifiable/observable to the principal, proper scoring rules are not applicable. However, the principal may nevertheless be interested in ensuring that agents are incentivized to report their beliefs truthfully. A leading example from Prelec (2004) is the case of subjective surveys/evaluations of a new product. In their example, the principal (e.g., a marketer) wishes to incentivize truthful reporting by agents, but there is no "ground truth" on which to base the rewards. The basic idea of Bayesian Truth Serum of Prelec (2004) is that agents' beliefs also have implications on their second-order beliefs (i.e., how they think the population will evaluate them). The population averages are indeed verifiable, and therefore can be used to incentivize agents. A similar idea works here, we describe it loosely for completeness:

(1) Elicit beliefs and second-order beliefs from agents as prescribed in PMBA to learn the true state ω (which variant of PMBA is used, and what is elicited, depends on the setting).
(2) For the agents who were asked for second-order beliefs, pay them based on a proper scoring rule as a function of their expected population average beliefs, and the true calculated population average beliefs.
(3) For all other agents, pay them using a proper scoring rule as a function of their beliefs and the state ω recovered by PMBA.

[17] There is a large literature proposing more robust scoring rules for various settings, see e.g., Karni (2009) for an example.

Observe that under the maintained assumptions of correctness of the corresponding PMBA procedure, truthful reporting will be a perfect Bayesian equilibrium among agents. To understand why, first note that since there is a large population, the average beliefs are effectively observed: even if a single agent misreports, this does not affect the average. Therefore, rewarding the two agents with a proper scoring rule based on their expected population average beliefs is feasible and will incentivize truthful reporting. These agents have no incentive to misreport their first-order beliefs because there is no associated payment. Note that given this, the true underlying state can be correctly calculated. A proper scoring rule against this recovered state ensures the truthful reporting of first-order beliefs for all the agents. Finally, there is no "compound" incentive for an agent who reports both first- and second-order beliefs to misreport their beliefs: their incentive to report their second-order beliefs truthfully was already argued, and misreporting their first-order beliefs will cause the PMBA to fail and therefore will void payments.

5. Bayesian Aggregation with Fixed Finite Populations

So far, we have discussed the possibility of aggregating information with an infinite or arbitrarily population. In practice, of course, populations are finite. The procedures given above will correctly identify the true state with high probability in large populations. For instance, in our baseline model, even with a large but finite population, the average belief in the population will concentrate around the expected belief conditional on the true state. As long as the expected beliefs conditional on state are sufficiently different across the states, an appropriately modified procedure will recover the true state of the world with high probability.

Nevertheless, at a theoretical level, one may wonder if, with a finite population, there is any value to eliciting higher-order beliefs, ignoring the difficulties of such elicitation in practice. To make this potentially valuable, we need a more exacting benchmark, since, as we argued above, PMBA will already aggregate information with appropriately high probability. The benchmark we use therefore is one of a "full information posterior," i.e., can we, without knowledge of the underlying information structure P, and eliciting solely agents beliefs (and higher-order beliefs), nevertheless reach the same posterior beliefs as an omniscient agent who knew P and directly observed all the agents' signals? Our previous results answered this in the positive and showed that for the case of an infinite population (under the maintained assumptions), the output of the PMBA could identify the degenerate belief on the true state.

In this section, we answer these questions primarily in the negative. First we show that with a finite population, and with elicitation of the entire hierarchy of beliefs, an agnostic procedure can learn the prior and signals of each agent. This result is essentially a straightforward corollary of the results of Mertens, Sorin, and Zamir (2015). However, elicitation of the entire hierarchy is obviously impractical in real-world applications.

By contrast, we show that, for elicitation up to any finite-order of beliefs, there is an identification problem: there exist information structures where the exact same finite hierarchy can be realized among agents in both states.

Consider a finite set of agents $\{1, 2, ..., N\}$ with a common prior P defined on a finite set $\Omega \times S$ (where $S = \times_{i=1}^{N} S_i$ is the set of their signal profiles). The aggregator does not know P but can ask the agents to report their higher-order beliefs. In this section, we show that the aggregator can effectively elicit the full information posterior, provided that they can ask the agents to report their entire hierarchy of beliefs and that each hierarchy of beliefs uniquely identifies a signal of an agent.

To this end, we recall the standard formulation of higher-order beliefs by Mertens and Zamir (1985). Denote by s_i^k the kth-order belief of agent i over Ω. For instance, $s_i^1 = \text{marg}_\Omega P(\cdot|s_i)$,

$$s_i^2\left(\omega, s_{-i}^1\right) = P\left(\left\{\left(\omega', s'_{-i}\right) : \omega' = \omega \text{ and } s'^1_{-i} = s^1_{-i}\right\} |s_i\right),$$

and so on. Denote by $\tilde{s}_i$ the hierarchy of beliefs induced by signal s_i, i.e., $\tilde{s}_i = \left(s_i^1, s_i^2, ...\right)$. Let $\tilde{S}_i$ be the set of hierarchies of beliefs induced from S_i and $\tilde{P}$ the distribution induced by P on $\Omega \times \tilde{S}$. Moreover, by Mertens and Zamir (1985), each $\tilde{s}_i$ induces a belief $\tilde{\pi}_i(\tilde{s}_i)$ over $\Omega \times \tilde{S}_{-i}$, where $\tilde{S}_{-i} = \times_{j \neq i} \tilde{S}_j$.

THEOREM 4. *Suppose that $\tilde{\mu}$ has a finite support and the agents report $\tilde{s} = (\tilde{s}_i)_i$. Then there exists a procedure which recovers the "pooled information" posterior on the states, i.e., $\tilde{P}(\cdot|\tilde{s})$.*

One might ask whether such a result can be obtained while only asking agents for their higher-order beliefs upto some finite order (i.e., as opposed to the full infinite hierarchy). It is possible to show that in general, without further assumptions, this is not the case: the result of Theorem 4 cannot be achieved if the aggregator only knows the reported beliefs up to order m. Our argument is adapted from the leading example of Lipman (2003). We defer the details to Appendix B.1.

6. RELATED LITERATURE AND CONCLUSIONS

In what follows, we discuss the related literature more closely (with the benefit of being able to explain some of the ideas more formally with our notation). Finally, we conclude.

The literature on the aggregation of information in a population is perhaps far too vast to cite comprehensively. Multiple strands attempt to understand, in a variety of settings, when and how existing institutions aggregate dispersed information in strategic settings. These include studies on information aggregation in voting, a vast literature with its roots in Condorcet, but more recently studied formally starting from Feddersen and Pesendorfer (1997); aggregation in auctions (see e.g., Milgrom and Weber (1982) and subsequent literature) and of course information aggregation in markets (see e.g., Grossman and Stiglitz (1980)). Ostrovsky (2009) proposes a market scoring rule mechanism that achieves information aggregation for some classes of securities in a dynamic market, but the payment rule there requires the state to be ex-post observable.

We limit attention here explicitly to work that studies the design/ possibility of aggregation procedures/ algorithms, or institutions that solely exist to aggregate information.

A key idea in the former space (indeed, in a sense an idea that is also at the heart of Prelec, Seung, and McCoy (2017)) is the Bayesian Truth Serum (Prelec (2004)— this is a procedure to truthfully elicit subjective information from agents by rewarding them as a function of others' reports. This paper pioneered the idea that agents' higher-order beliefs could be used to incentivize them successfully even if the planner/ designer did not know the prior among agents. A literature studying this followed. Prelec and Seung (2006) showed how to modify agents' reports to aggregate them better (akin to the prediction normalized voting rule we described in Example 1). Several extensions to the BTS have been proposed— see e.g., Cvitanić, Prelec, Riley, and Tereick (2019), Witkowski and Parkes (2012), Radanovic and Faltings (2013), and Radanovic and Faltings (2014). The latter three papers attempt to make the basic Bayesian Truth Serum robust, i.e., to ensure that it has good properties even in finite populations and to ensure that participation is individually rational, etc. More recent work has also expanded the scope of the question to continuous spaces, but with more structure—see, e.g., Kong, Schoenebeck, Tao, and Yu (2020).

These ideas have also been successfully employed in real-world applications: see e.g., Shaw, Horton, and Chen (2011) for an application on Mechanical Turk, or Rigol, Hussam, and Roth (2021), who use the Robust BTS of Witkowski and Parkes (2012) to identify high-ability micro-entrepreneurs by surveying peers. There is also a growing literature attempting to elicit and use higher-order beliefs to improve aggregation—see e.g., recent works such as Palley and Soll (2019), Palley and Satopää (2020). Some of these works consider modifications of the SP algorithm that improve performance in experiments, e.g., Wilkening, Martinie, and Howe (2021) or Martinie, Wilkening, and Howe (2020). For example, Wilkening, Martinie, and Howe (2021) show that opposite to SP, SC weights forecasters with more informative private signals more than forecasters with less informative ones. This explains why SC is favorable relative to SP when both guess/vote and posterior data are available. Similarly, McCoy and Prelec (0) use Bayesian Hierarchical methods to improve performance relative to SP.

On the latter, there is an important literature which has attempted to understand the performance of prediction markets: see e.g., Wolfers and Zitzewitz (2004) for an early piece summarizing the issues and Baillon (2017) for an alternate design. Wolfers and Zitzewitz (2006) study when the prediction market price can be interpreted as the average of traders' beliefs, while Ottaviani and Sørensen (2015) study when the market price under-reacts to new information. Dai, Jia, and Kou (2020) show how to infer the state of the world from prediction market trading data using these theoretical ideas. There is also a large literature on pari-mutuel markets, and on adapting them to aggregate information—see e.g., Pennock (2004).

As we mentioned in footnote 3, our approach has some similiarities with Theorem 1.4 of Prelec, Seung, and McCoy (2017). This approach require elicitation of $p(s_i|s_j)$, i.e. the conditional probability that a different agent has seen signal s_i given that you have seen s_j. In particular this requires that the subjects understand the meaning of signals (H and T in their example). In contrast, PMBA questions are all in terms of primitives (beliefs and higher-order beliefs) of an agent. Our position is that this distinction is important because if the subjects really understand the model/prior including meaning of signals, the easiest question is to simply ask them for the state-contingent population signal distribution (i.e. the bias of the various possible coins in the terminology of PSM).

A different approach to ours for aggregating information with more than 2 states can be found in Libgober (2021). The principal does not know the agent's prior, but can elicit, after an unknown Blackwell experiment, their posterior and their contingent hypothetical beliefs (their posteriors contingent on seeing any signal). He shows that the prior and experiment can be recovered from these contingent hypothetical beliefs, via a linear-regression method, even if some signals are not sampled (i.e. μ need not be a square matrix). In contrast, Prelec and McCoy (2022), following Samet (1998b), propose to identify contingent hypothetical beliefs with a transition matrix on signals. An invariant distribution of the matrix corresponds to the ex-ante population signal distribution and thereby delivers the unknown the information structure. At a technical level, our "expected population average beliefs" serve a similar role to the "contingent hypothetical beliefs" in these papers. However, the contingent hypothetical beliefs include information about higher-order beliefs well beyond the expected population mean; moreover, both Libgober (2021) and Prelec and McCoy (2022) require that the agents share a common prior and thereby rule out misspecified information settings (Procedure 3).

On a related theme, there are the works of Samet (1998a), Samet (1998b), and Golub and Morris (2017). These papers are interested in interim characterizations of the common prior, i.e., understanding whether or not agents have a common prior from the properties of their interim beliefs/higher-order beliefs.

References

ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): "Bayesian learning in social networks," *The Review of Economic Studies*, 78(4), 1201–1236.

ARIELI, I., Y. BABICHENKO, AND R. SMORODINSKY (2017): "When is the Crowd Wise?," *Available at SSRN 3083608*.

AUSTEN-SMITH, D., AND J. S. BANKS (1996): "Information aggregation, rationality, and the Condorcet jury theorem," *American political science review*, 90(1), 34–45.

BAILLON, A. (2017): "Bayesian markets to elicit private information," *Proceedings of the National Academy of Sciences*, 114(30), 7958–7962.

BOHREN, J. A. (2016): "Informational herding with model misspecification," *Journal of Economic Theory*, 163, 222–247.

BRIER, G. W. (1950): "Verification of forecasts expressed in terms of probability," *Monthly weather review*, 78(1), 1–3.

CHARNESS, G., AND D. LEVIN (2009): "The Origin of the Winner's Curse: A Laboratory Study," *American Economic Journal: Microeconomics*, 1(1), 207–36.

CVITANIĆ, J., D. PRELEC, B. RILEY, AND B. TEREICK (2019): "Honesty via choice-matching," *American Economic Review: Insights*, 1(2), 179–92.

DAI, M., Y. JIA, AND S. KOU (2020): "The wisdom of the crowd and prediction markets," *Journal of Econometrics*.

ESPONDA, I., AND E. VESPA (2014): "Hypothetical Thinking and Information Extraction in the Laboratory," *American Economic Journal: Microeconomics*, 6(4), 180–202.

EYSTER, E., AND M. RABIN (2010): "Naive herding in rich-information settings," *American economic journal: microeconomics*, 2(4), 221–43.

FEDDERSEN, T., AND W. PESENDORFER (1997): "Voting behavior and information aggregation in elections with private information," *Econometrica: Journal of the Econometric Society*, pp. 1029–1058.

FRICK, M., R. IIJIMA, AND Y. ISHII (2020): "Misinterpreting others and the fragility of social learning," *Econometrica*, 88(6), 2281–2328.

FU, H., N. HAGHPANAH, J. HARTLINE, AND R. KLEINBERG (2021): "Full surplus extraction from samples," *Journal of Economic Theory*, 193, 105230.

GOLUB, B., AND M. O. JACKSON (2010): "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, 2(1), 112–49.

GOLUB, B., AND S. MORRIS (2017): "Higher-Order Expectations," *Available at SSRN 2979089*.

GROSSMAN, S. J., AND J. E. STIGLITZ (1980): "On the impossibility of informationally efficient markets," *The American economic review*, 70(3), 393–408.

GUL, F. (1998): "A comment on Aumann's Bayesian view," *Econometrica*, 66(4), 923–927.

KARNI, E. (2009): "A mechanism for eliciting probabilities," *Econometrica*, 77(2), 603–606.

KONG, Y., G. SCHOENEBECK, B. TAO, AND F.-Y. YU (2020): "Information elicitation mechanisms for statistical estimation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 2095–2102.

LI, S. (2017): "Obviously Strategy-Proof Mechanisms," *American Economic Review*, 107(11), 3257–87.

LIBGOBER, J. (2021): "Hypothetical Beliefs Identify Information," .

LIPMAN, B. L. (2003): "Finite order implications of common priors," *Econometrica*, 71(4), 1255–1267.

MARTINIE, M., T. WILKENING, AND P. D. L. HOWE (2020): "Using meta-predictions to identify experts in the crowd when past performance is unknown," *PLOS ONE*, 15(4), 1–11.

MARTÍNEZ-MARQUINA, A., M. NIEDERLE, AND E. VESPA (2019): "Failures in Contingent Reasoning: The Role of Uncertainty," *American Economic Review*, 109(10), 3437–74.

MCCOY, J., AND D. PRELEC (0): "A Bayesian Hierarchical Model of Crowd Wisdom Based on Predicting Opinions of Others," *Management Science*, 0(0), null.

MERTENS, J.-F., S. SORIN, AND S. ZAMIR (2015): *Repeated games*, vol. 55. Cambridge University Press.

MERTENS, J.-F., AND S. ZAMIR (1985): "Formulation of Bayesian analysis for games with incomplete information," *International Journal of Game Theory*, 14(1), 1–29.

MILGROM, P. R., AND R. J. WEBER (1982): "A theory of auctions and competitive bidding," *Econometrica: Journal of the Econometric Society*, pp. 1089–1122.

NGANGOUÉ, M. K., AND G. WEIZSÄCKER (2021): "Learning from Unrealized versus Realized Prices," *American Economic Journal: Microeconomics*, 13(2), 174–201.

OSTROVSKY, M. (2009): "Information aggregation in dynamic markets with strategic traders," in *Proceedings of the 10th ACM conference on Electronic commerce*, pp. 253–254.

OTTAVIANI, M., AND P. N. SØRENSEN (2015): "Price reaction to information with heterogeneous beliefs and wealth effects: Underreaction, momentum, and reversal," *American Economic Review*, 105(1), 1–34.

PALLEY, A., AND V. SATOPÄÄ (2020): "Boosting the wisdom of crowds within a single judgment problem: Selective averaging based on peer predictions," *Available at SSRN 3504286*.

PALLEY, A. B., AND J. B. SOLL (2019): "Extracting the wisdom of crowds when information is shared," *Management Science*, 65(5), 2291–2309.

PEKER, C., AND T. WILKENING (2023): "Robust recalibration of aggregate probability forecasts using meta-beliefs," .

PENNOCK, D. M. (2004): "A dynamic pari-mutuel market for hedging, wagering, and information aggregation," in *Proceedings of the 5th ACM conference on Electronic commerce*, pp. 170–179.

PRELEC, D. (2004): "A Bayesian truth serum for subjective data," *science*, 306(5695), 462–466.

PRELEC, D., AND J. MCCOY (2022): "General identifiability of possible world models for crowd wisdom," .

PRELEC, D., H. S. SEUNG, AND J. MCCOY (2017): "A solution to the single-question crowd wisdom problem," *Nature*, 541(7638), 532.

PRELEC, D., AND S. SEUNG (2006): "An algorithm that finds truth even if most people are wrong," *Unpublished manuscript*, 69.

RADANOVIC, G., AND B. FALTINGS (2013): "A robust bayesian truth serum for non-binary signals," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 27.

——— (2014): "Incentives for truthful information elicitation of continuous signals," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 28.

RIGOL, N., R. HUSSAM, AND B. ROTH (2021): "Targeting high ability entrepreneurs using community information: Mechanism design in the field," *American Economic Review, forthcoming*.

SAMET, D. (1998a): "Common priors and separation of convex sets," *Games and economic behavior*, 24(1-2), 172–174.

——— (1998b): "Iterated expectations and common priors," *Games and economic Behavior*, 24(1-2), 131–141.

SHAW, A. D., J. J. HORTON, AND D. L. CHEN (2011): "Designing incentives for inexpert human raters," in *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, pp. 275–284.

SHIRYAEV, A. N. (2012): *Problems in probability*. Springer Science & Business Media.

WILKENING, T., M. MARTINIE, AND P. D. L. HOWE (2021): "Hidden Experts in the Crowd: Using Meta-Predictions to Leverage Expertise in Single-Question Prediction Problems," *Management Science, Forthcoming*.

WITKOWSKI, J., AND D. PARKES (2012): "A robust bayesian truth serum for small populations," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 26.

WOLFERS, J., AND E. ZITZEWITZ (2004): "Prediction markets," *Journal of economic perspectives*, 18(2), 107–126.

——— (2006): "Prediction markets in theory and practice," Discussion paper, national bureau of economic research.

APPENDIX A. PROOFS FROM SECTION 2

First we prove a general property:

LEMMA 1. *Suppose Part (2) of Assumption 1 is satisfied. Then for any state ω,*

$$\lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{\infty} \tilde{\mu}_i \to_P \overline{\mu}(\omega).$$

PROOF. The result is established in two steps. For any state ω below, we show that agents' beliefs that the state is $\omega' \in \Omega$ satisfy vanishing correlation. As beliefs lie in $[0,1]$ and therefore have bounded variance, we can apply the Bernstein law of large numbers. Together, these imply that when the state of the world is ω; for any state $\omega' \in \Omega$, we have that:

$$\lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{\infty} \tilde{\mu}_{i,\omega'} \to_P \overline{\mu}_{\omega'}(\omega).$$

Since there are only a finite set of states, the main result follows.

First, let us prove our claim of vanishing correlations:

CLAIM 2. *Conditioned on state $\omega \in \Omega$, the sequence of (random) beliefs $\langle \tilde{\mu}_{i,\omega'} \rangle_{i=1}^{\infty}$ that the state of the world is any $\omega' \in \Omega$ satisfies vanishing correlation, i.e.,*

$$\lim_{|i-j|\to\infty} cov(\tilde{\mu}_{i,\omega'}, \tilde{\mu}_{j,\omega'}|\omega) = 0.$$

PROOF OF CLAIM. Consider agents i and j such that ε-independence among signals holds, i.e., for all events $E, E' \in \mathcal{F}$ we have

$$|P\left[\tilde{s}_i \in E, \tilde{s}_j \in E'|\omega\right] - P^S_{i,\omega}(E)P^S_{j,\omega}(E')| \leq \varepsilon.$$

Let f_i denote the conditional probability function that assigns Bayesian beliefs of state ω' to signals for agent i. By measurability of f_i's we have that every event $S, S' \subseteq [0,1]$ has a corresponding event $f_i^{-1}(S)$ and $f_j^{-1}(S')$ such that

$$\begin{aligned}&\left|\Pr\left[\tilde{\mu}_{i,\omega'} \in S, \tilde{\mu}_{j,\omega'} \in S' \,|\omega\right] - \Pr\left[\tilde{\mu}_{i,\omega'} \in S\,|\omega\right]\Pr\left[\tilde{\mu}_{j,\omega'} \in S'\,|\omega\right]\right| \\ =&\left|P\left[\tilde{s}_i \in f_i^{-1}(S), \tilde{s}_j \in f_j^{-1}(S')\,|\omega\right] - P^S_{i,\omega}(f_i^{-1}(S))P^S_{j,\omega}(f_j^{-1}(S'))\right|\end{aligned}$$

establishing approximate independence of beliefs, conditional on the state.

To establish state dependent vanishing correlation among beliefs, consider the random beliefs $\tilde{\mu}_{i,\omega'}, \tilde{\mu}_{j,\omega'}$. For any number $n > 1$, let G_n denote a partition of $[0,1]$ into $n+1$ equal intervals, i.e., $G_n^k = [\frac{k-1}{n}, \frac{k}{n})$ for $k = 1$ thru $n-1$; $G_n^n = [\frac{n-1}{n}, 1]$.

Note that:

$$\mathbb{E}[\tilde{\mu}_{i,\omega'}|\omega] \leq \sum_{k=1}^{n} \frac{k}{n}\Pr(\tilde{\mu}_{i,\omega'} \in G_n^k|\omega) \equiv U_i^n,$$

$$\mathbb{E}[\tilde{\mu}_{j,\omega'}|\omega] \leq \sum_{k=1}^{n} \frac{k}{n}\Pr(\tilde{\mu}_{j,\omega'} \in G_n^k|\omega) \equiv U_j^n,$$

$$\mathbb{E}[\tilde{\mu}_{i,\omega'}\tilde{\mu}_{j,\omega'}|\omega] \leq \sum_{l=1}^{n}\sum_{k=1}^{n}\frac{lk}{n}\Pr(\tilde{\mu}_{i,\omega'} \in G_n^l, \tilde{\mu}_{j,\omega'} \in G_n^k|\omega) \equiv U_{ij}^n.$$

Since $\tilde{\mu}_{i,\omega'}$ and $\tilde{\mu}_{j,\tilde{\mu}'}$ satisfy ε-independence, we have that

$$|U_{ij}^n - U_i^n U_j^n| \leq \varepsilon.$$

However, as n goes to infinity, by observation, $U_i^n \to \mathbb{E}[\tilde{\mu}_{i,\omega'}|\omega]$, $U_j^n \to \mathbb{E}[\tilde{\mu}_{j,\omega'}|\omega]$ and $U_{ij}^n \to \mathbb{E}[\tilde{\mu}_{i,\omega'}\tilde{\mu}_{j,\omega'}|\omega]$. Therefore $\text{cov}(\tilde{\mu}_{i,\omega'}, \tilde{\mu}_{j,\omega'}|\omega) \leq \varepsilon$.

Now, by part (2) of Assumption 1, we know that for any i, and $\varepsilon > 0$, $\exists n(\varepsilon)$ large enough so that if $j - i > n$ so that we have that conditional on state, $\tilde{\mu}_i$ and $\tilde{\mu}_j$ are ε-independent. The result trivially follows. ■

Now, recall the Bernstein law of large numbers (See e.g 3.3.2 of Shiryaev (2012)):

THEOREM 5. *Let $\tilde{X}_i$, $i = 1,2,3\ldots$ be a sequence of real valued random variables such that (1) $\mathbb{E}[\tilde{X}_i]$ is finite for each i, (2) there exists K finite such that $\mathbb{V}[\tilde{X}_i] \leq K$ for all i, and (3) $\lim_{|i-j|\to\infty} cov(\tilde{X}_i, \tilde{X}_j) = 0$. Then*

$$\lim_{n\to\infty}\frac{1}{n}\sum_{i=1}^{n}\left(\tilde{X}_i - \mathbb{E}[\tilde{X}_i]\right) \to_P 0.$$

Observe that since $\tilde{\mu}_{i,\omega'}$ is a belief, it lies between $[0,1]$ and so the first two conditions of the Theorem are verified. Claim 2 verifies the third condition conditional on any true state ω. Therefore we have the desired result for any $\omega' \in \Omega$, conditional on true state ω. Since the set of states is finite, we have convergence in probability of the entire vector. ■

LEMMA 2. *Suppose $L = 2$. Suppose further that part (1) of Assumption 1 holds and $P_{i,\omega_1}^S \neq P_{i,\omega_2}^S$. Then following statements hold:*

(1) $\frac{\hat{G}_{i,\omega_1}(r)}{\hat{G}_{i,\omega_2}(r)}$ is non-increasing in r

(2) $\frac{\hat{G}_{i,\omega_1}(r)}{\hat{G}_{i,\omega_2}(r)} > 1$ for all $r \in \left(\beta, \bar{\beta}\right)$ where $\left[\underline{\beta}, \bar{\beta}\right]$ denote the convex hull of the support of private beliefs.[18]

(3) $\overline{\mu}_i(\omega_1) < \overline{\mu}_i(\omega_2)$.

PROOF. By Bayes' rule, we know that

$$r = \frac{dG_{i,\omega_2}(r)\,P(\omega_2)}{dG_{i,\omega_1}(r)\,P(\omega_1) + dG_{i,\omega_2}(r)\,P(\omega_2)}. \tag{7}$$

[18] Since signal generating measures are absolutely continuous each other, the support is identical across states.

We define $\alpha = P(\omega_2)/P(\omega_1)$ as the relative likelihood of the states according to the prior P. Since no signal is completely informative, i.e., $r \notin \{0,1\}$, we can rewrite (7) as

$$\frac{dG_{i,\omega_1}}{dG_{i,\omega_2}}(r) = \left(\frac{1-r}{r}\right)\alpha. \tag{8}$$

Hence, for any $r \in (0,1)$, we have

$$G_{i,\omega_1}(r) = \int_0^r dG_{i,\omega_1}(x) = \int_0^r \left(\frac{1-x}{x}\right)\alpha dG_{i,\omega_2}(x) \geq \left(\frac{1-r}{r}\right)\alpha G_{i,\omega_2}(r). \tag{9}$$

Part (1). Pick any r', r s.t. $r' > r$.Observe that:

$$\begin{aligned}
&\frac{G_{i,\omega_1}(r)}{G_{i,\omega_2}(r)} - \frac{G_{i,\omega_1}(r')}{G_{i,\omega_2}(r')},\\
=&\frac{G_{i,\omega_1}(r)G_{i,\omega_2}(r') - G_{i,\omega_1}(r')G_{i,\omega_2}(r)}{G_{i,\omega_2}(r)G_{i,\omega_2}(r')}.
\end{aligned}$$

Define $\Delta_{i,\omega} = G_{i,\omega}(r') - G_{i,\omega}(r)$. We then have:

$$=\frac{G_{i,\omega_1}(r)\Delta_{i,\omega_2} - \Delta_{i,\omega_1}G_{i,\omega_2}(r)}{G_{i,\omega_2}(r)G_{i,\omega_2}(r')},$$

Observe that $\Delta_{i,\omega_1} = \int_r^{r'} dG_{i,\omega_1}(x) = \int_r^{r'}\left(\frac{1-x}{x}\right)\alpha dG_{i,\omega_2}(x) \leq \left(\frac{1-r}{r}\right)\alpha\Delta_{i,\omega_2}$. Substituting in,

$$\begin{aligned}
\geq&\frac{G_{i,\omega_1}(r)\Delta_{i,\omega_2} - \frac{1-r}{r}\alpha\Delta_{i,\omega_2}G_{i,\omega_2}(r)}{G_{i,\omega_2}(r)G_{i,\omega_2}(r')},\\
\geq&0,
\end{aligned}$$

where the last inequality follows by substituting in $G_{i,\omega_1}(r)$ in the numerator from (9).
Part (2). It follows from (9) that

$$G_{i,\omega_1}(r) \geq G_{i,\omega_2}(r) \text{ if } \left(\frac{1-r}{r}\right)\alpha \geq 1. \tag{10}$$

If $\left(\frac{1-r}{r}\right)\alpha < 1$, then

$$\begin{aligned}
1 - G_{i,\omega_1}(r) &= \int_r^1 dG_{i,\omega_1}(x)\\
&= \int_r^1 \left(\frac{1-x}{x}\right)\alpha dG_{i,\omega_2}(x)\\
&\leq \left(\frac{1-r}{r}\right)\alpha\left(1 - G_{i,\omega_2}(r)\right)\\
&\leq 1 - G_{i,\omega_2}(r).
\end{aligned} \tag{11}$$

Combining (10) and (11), we obtain that

$$\frac{G_{i,\omega_1}(r)}{G_{i,\omega_2}(r)} \geq 1. \tag{12}$$

To prove $\frac{G_{i,\omega_1}(r)}{G_{i,\omega_2}(r)} > 1$ for any $r \in \left(\underline{\beta}, \overline{\beta}\right)$, suppose to the contrary that $\frac{G_{i,\omega_1}(r)}{G_{i,\omega_2}(r)} = 1$ for some $r \in \left(\underline{\beta}, \overline{\beta}\right)$. We derive contradiction in each of the following two cases

Case 1: $\left(\frac{1-r}{r}\right)\alpha < 1$. Then,

$$\begin{aligned} G_{i,\omega_1}(1) &= G_{i,\omega_1}(r) + \int_r^1 dG_{i,\omega_1}(x) \\ &= G_{i,\omega_2}(r) + \int_r^1 dG_{i,\omega_1}(x) \\ &= G_{i,\omega_2}(r) + \int_r^1 \left(\frac{1-x}{x}\right)\alpha dG_{i,\omega_2}(x) \\ &\leq G_{i,\omega_2}(r) + \left(\frac{1-r}{r}\right)\alpha \int_r^1 dG_{i,\omega_2}(x) \\ &= G_{i,\omega_2}(r) + \left(\frac{1-r}{r}\right)\alpha \left[1 - G_{i,\omega_2}(r)\right] \end{aligned}$$

which yields a contradiction unless $G_{i,\omega_2}(r) = 1$. However, $G_{i,\omega_2}(r) = 1$ implies $r \geq \overline{\beta}$ which contradicts $r \in \left(\underline{\beta}, \overline{\beta}\right)$.

Case 2: $\left(\frac{1-r}{r}\right)\alpha \geq 1$. Then, since $\frac{G_{i,\omega_1}(r)}{G_{i,\omega_2}(r)} = 1$ and $\frac{G_{i,\omega_1}(r)}{G_{i,\omega_2}(r)}$ is non-increasing in r, together with (12) implies that $\frac{G_{i,\omega_1}(r')}{G_{i,\omega_2}(r')} = 1$ for all $r' \geq r$. Pick $r' > \frac{\alpha}{1+\alpha}$ (note that since $\frac{\alpha}{1+\alpha}$ is just the prior belief that $\omega_1 = 1$, it must be in the interior of the support of beliefs $[\underline{\beta}, \bar{\beta}]$—this leads to the contradiction in Case 1.

Part (3). This follows directly from the first-order stochastic dominance relationship between G_{i,ω_1} and G_{i,ω_2} in established part (2). ■

LEMMA 3. *Suppose $L = 2$ and the common prior P satisfies both Assumptions 1 and 2. Then, there exists an $\varepsilon > 0$ such that the $\overline{\mu}_i(\omega_1) + \varepsilon < \overline{\mu}_i(\omega_2)$ for all $i \in N$.*

PROOF. Fix any agent i and let the convex hull of the common support of its posteriors be $[\underline{\beta}, \overline{\beta}]$. First observe that the prior belief about the state is $\frac{\alpha}{1+\alpha}$. Therefore it must be that $\underline{\beta} < \frac{\alpha}{1+\alpha} < \overline{\beta}$ (if exactly one of these is an equality we violate the condition that the expected posterior is the prior, if both are equalities then the signal is uninformative violating Assumption 2.

CLAIM 3. *For $\delta > 0$ if the total variation distance between G_{i,ω_1} and G_{i,ω_2} is larger than δ, there exists $\delta' > 0$ such that $\underline{\beta} + \delta' < \frac{\alpha}{1+\alpha} < \overline{\beta} - \delta'$, and further*

$$\max\left\{G_{i,\omega_1}(\frac{\alpha}{1+\alpha} - \delta'), (1 - G_{i,\omega_2}(\frac{\alpha}{1+\alpha} + \delta')\right\} \geq \delta'.$$

PROOF. To see this observe that we showed in the proof of Lemma 2 that for any $r \notin \{0,1\}$, we have

$$\frac{dG_{i,\omega_1}}{dG_{i,\omega_2}}(r) = \left(\frac{1-r}{r}\right)\alpha.$$

Therefore for r arbitrarily close to $\frac{\alpha}{1+\alpha}$, $\frac{dG_{i,\omega_1}}{dG_{i,\omega_2}}(r)$ is arbitrarily close to 1. Further, by the fact that G_{i,ω_2} first order stochastically dominates G_{i,ω_1}, we have $\delta' > G_{i,\omega_1}(\frac{\alpha}{1+\alpha} - \delta') > G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta')$, and similarly $\delta' > (1 - G_{i,\omega_2}(\frac{\alpha}{1+\alpha} + \delta') > (1 - G_{i,\omega_1}(\frac{\alpha}{1+\alpha} + \delta')$

Therefore if the claim is not satisfied, both measures are mostly supported on $[\frac{\alpha}{1+\alpha} - \delta', \frac{\alpha}{1+\alpha} + \delta']$ the total variation distance is $\leq 2\frac{1+\alpha}{\alpha}\delta'^2 + 2\delta'$. Therefore the claim follows for e.g., $\delta' = \frac{1}{3}\delta$. ∎

Now observe that $\overline{\mu}_i(\omega) = \int_0^1 (1 - G_{i,\omega}(r))dr$. Therefore

$$\begin{aligned}&\overline{\mu}_i(\omega_2) - \overline{\mu}(\omega_1)\\ =&\int_0^1 (G_{i,\omega_1}(r) - G_{i,\omega_2}(r))dr\end{aligned}$$

We showed in Lemma 2 that G_{i,ω_2} first order stochastically dominates G_{i,ω_1}. Therefore we have, for δ' that satisfies the statement of the claim above,

$$\begin{aligned}\geq& \int_{\frac{\alpha}{1+\alpha}-\delta'}^{\frac{\alpha}{1+\alpha}} (G_{i,\omega_1}(r) - G_{i,\omega_2}(r))dr\\ =&\delta'(G_{i,\omega_1}(\frac{\alpha}{1+\alpha} - \delta') - G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta'))\\ &+ \int_{\frac{\alpha}{1+\alpha}-\delta'}^{\frac{\alpha}{1+\alpha}} \left(\int_{\frac{\alpha}{1+\alpha}-\delta'}^{r} (dG_{i,\omega_1}(x) - dG_{i,\omega_2}(x))\right) dr\end{aligned}$$

We know from (8) that $dG_{i,\omega_1}(r) > dG_{i,\omega_2}(r)$ for $r < \frac{\alpha}{1+\alpha}$. So we have,

$$\geq \delta'(G_{i,\omega_1}(\frac{\alpha}{1+\alpha} - \delta') - G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta')).$$

Now observe that by (9) we have

$$\begin{aligned}&G_{i,\omega_1}(r) \geq \frac{1-r}{r}\alpha G_{i,\omega_2}(r),\\ \implies& G_{i,\omega_1}(r) - G_{i,\omega_2}(r) \geq \left(\frac{1-r}{r}\alpha - 1\right) G_{i,\omega_2}(r).\end{aligned}$$

For $r = \frac{\alpha}{1+\alpha} - \delta'$ we have

$$\Longrightarrow G_{i,\omega_1}(\frac{\alpha}{1+\alpha} - \delta') - G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta') \geq \frac{\delta'(\alpha+1)^2}{\alpha - (\alpha+1)\delta'} G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta').$$

However we also have by the statement of the claim that

$$G_{i,\omega_1}(\frac{\alpha}{1+\alpha} - \delta') \geq \delta',$$
$$\Longrightarrow G_{i,\omega_1}(\frac{\alpha}{1+\alpha} - \delta') - G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta') \geq \delta' - G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta').$$

Combining, we have that

$$\geq \max\{\delta' - G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta'), \frac{\delta'(\alpha+1)^2}{\alpha - (\alpha+1)\delta'} G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta')\}$$

Note that the first term is decreasing in $G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta')$, while the second is increasing. Therefore, the minimum is achieved at

$$\delta' - G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta') = \frac{\delta'(\alpha+1)^2}{\alpha - (\alpha+1)\delta'} G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta')$$
$$\Longrightarrow G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta') = \frac{\alpha - (\alpha+1)\delta'}{\alpha + \alpha\delta'(\alpha+1)}\delta'$$

Substituting in we have that,

$$G_{i,\omega_1}(\frac{\alpha}{1+\alpha} - \delta') - G_{i,\omega_2}(\frac{\alpha}{1+\alpha} - \delta') \geq \frac{\delta'(\alpha+1)^2}{\alpha + \alpha\delta'(\alpha+1)}\delta',$$

and therefore we have that $\overline{\mu}_i(\omega_2) - \overline{\mu}_i(\omega_1) \geq \varepsilon$ for $\varepsilon = \frac{\delta'(\alpha+1)^2}{\alpha+\alpha\delta'(\alpha+1)}\delta'^2$. ∎

We are now finally in a position to give the proof of Theorem 1.

THEOREM 1. *Suppose the information structure among agents satisfies Assumptions 1 and 2. Then the PMBA procedure (Procedure 1) correctly recovers the true state of the world almost surely.*

PROOF OF THEOREM 1. We prove this by analyzing each of the steps of Procedure 1.

(1) By Lemma 1, the law of large numbers holds and when the realizes state is ω,
$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \tilde{\mu}_i = \overline{\mu}(\omega).$$
Therefore the population mean of the elicited beliefs $\hat{\mu}$ in Step (1) of the procedure equals $\overline{\mu}(\omega)$ in state ω.

(2) By part (1) of Assumption 1, we know that the beliefs of any agent i have a common support across the states. Further, by Assumption 2, this support must have

at least 2 distinct points in it. Combining this with the fact that signals are approximately independent, there must exist at least two agents with different beliefs almost surely. So step 2 can find 2 such agents. Further, in step 2, each agent $i = A, B$ will report, by the law of total probability: $\alpha_i = \mu_i \overline{\mu}(\omega_1) + (1 - \mu_i)\overline{\mu}(\omega_2)$.

(3) Since $\mu_A \neq \mu_B$, equation 1 can be solved.

(4) By Lemma 3, there is a $\varepsilon > 0$ such that for all i $\overline{\mu}_i(\omega_1) < \overline{\mu}_i(\omega_2) - \varepsilon$. Therefore, the expected poptulation average beliefs, $\overline{\mu}(\omega_1) < \overline{\mu}(\omega_2) - \varepsilon$, i.e., $\overline{\mu}(\omega_1) \neq \overline{\mu}(\omega_2)$. Therefore comparing the elicited $\hat{\mu}$ in step (1) with the recovered $(\overline{\mu}(\omega_1), \overline{\mu}(\omega_2))$ must reveal the true state. ∎

THEOREM 2. *Suppose $L > 2$ and the common prior P satisfies Assumptions 1, 3, and 4. Then the PMBA procedure (Procedure 2) recovers the true unknown state of nature ω almost surely.*

PROOF. We prove this by analyzing each of the steps of Procedure 2 analogously to our proof of Theorem 1.

(1) By Lemma 1, the law of large numbers holds and when the realizes state is ω,

$$\lim_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} \tilde{\mu}_i = \overline{\mu}(\omega).$$

Therefore the population mean of the elicited beliefs $\hat{\mu}$ in Step (1) of the procedure equals $\overline{\mu}(\omega)$ in state ω.

(2) By part (1) of Assumption 1, we know that the beliefs of any agent i have a common support across the states. Further, by Assumption 4, this support must has an interior relative to Δ^L. Combining this with the fact that signals are approximately independent, there must exist at least L agents whose beliefs form a full rank matrix a.s.. So step 2 can find L such agents. Further, in step 2, each agent $i \in I$ will report, by the law of total probability: $\alpha_i = \mu_i^T \bar{\boldsymbol{\mu}}$.

(3) Since the set I of agents is selected so that $\boldsymbol{\mu}$ is full-rank, equation 2 can be solved.

(4) By Assumption 3, $\overline{\mu}(\omega) \neq \overline{\mu}(\omega')$ for any $\omega, \omega' \in \Omega$.

Therefore comparing the elicited $\hat{\mu}$ in step (1) with the recovered $\bar{\boldsymbol{\mu}}$ must reveal the true state. ∎

CLAIM 1. *Suppose that signals are conditionally i.i.d. draws from a finite set S, and Assumption 3 is satisfied. If each agent gets $L - 1$ independent draws from the distribution P^S_ω, then Assumption 4 is satisfied with the signal space S^{L-1}.*

PROOF. At each ω, the distribution of the $L - 1$ signals can be summarized by the outer product

$$\left(\otimes P^S_\omega\right)^{L-1} = \underbrace{P^S_\omega \otimes \cdots \otimes P^S_\omega}_{L-1 \text{ times}}.$$

Since Assumption 3 is satisfied, the dimension of the linear space spanned by $\left\{P_{\omega}^{S}\right\}_{\omega\in\Omega}$ is at least 2. Hence, by Lemma 4 by Fu, Haghpanah, Hartline, and Kleinberg (2021), $\left\{\left(\otimes P_{\omega}^{S}\right)^{L-1}\right\}_{\omega\in\Omega}$ are linearly independent and hence $\left\{\left(\otimes P_{\omega}^{S}\right)^{L-1}P\left(\omega\right)\right\}_{\omega\in\Omega}$ are also linearly independent ($P\left(\omega\right)$ denotes the prior probability of ω). Identify $\left\{\left(\otimes P_{\omega}^{S}\right)^{L-1}P\left(\omega\right)\right\}_{\omega\in\Omega}$ with a $|S|^{L-1}\times L$ matrix which we denote by M. Hence, rank$(M)=L$. For $\mathbf{s}=\left(s^{1},...,s^{L-1}\right)$ and $\omega\in\Omega$, observe that $M_{\mathbf{s},\omega}=P\left(\mathbf{s},\omega\right)$. Since rank$(M)=L$, the conditional probability matrix $\left\{P\left(\omega|\mathbf{s}\right)\right\}$ also has rank L. Hence, Assumption 4 is satisfied with the signal space S^{L-1}. ■

A.1. *Proof of Theorem 3*

We first need some more lemmata.

LEMMA 4. *Let $\tilde{\mu}_A$,$\tilde{\mu}_B$ denote the random group mean beliefs and μ_A,μ_B their expectation. We have that $\tilde{\mu}_A$(respectively, $\tilde{\mu}_B$) converges in probability to the expectation μ_A(respectively, μ_B).*

PROOF. Let x_{iA},x_{iB} denote the indicator random variable of i belonging to group A,B, and let μ_i denote the random belief of agent i. We have that

$$\tilde{\mu}_A=\lim_{n\in\mathbb{N}}\frac{1}{\sum\limits_{i<n}x_{iA}}\sum_{i<n}x_{iA}\mu_i\ =\lim_{n\in\mathbb{N}}\frac{1}{\frac{1}{n}\sum\limits_{i<n}x_{iA}}\frac{1}{n}\sum_{i<n}x_{iA}\mu_i.$$

Let

$$L_n=\frac{1}{\frac{1}{n}\sum\limits_{i<n}x_{iA}},\text{ and, }K_n=\frac{1}{n}\sum_{i<n}x_{iA}\mu_i.$$

Note that since the conditions for Bernstein's LLN are satisfied for the sequences $\langle x_{iA}\rangle$ and $\langle x_{iA}\mu_i\rangle$ conditional on the realized state, it follows that $\langle L_n\rangle$ and $\langle K_n\rangle$ converge in probability to L respectively K where

$$L=\lim_{n\in\mathbb{N}}E[L_n],\ K=\lim_{n\in\mathbb{N}}E[K_n].$$

Since $\langle L_n\rangle$ and $\langle K_n\rangle$ are uniformly bounded, it follows that $\tilde{\mu}_A$ converges in probability to μ_A

$$\mu_A=\lim_{n\in\mathbb{N}}\frac{1}{E\left[\frac{1}{n}\sum\limits_{i<n}x_{iA}\right]}E\left[\frac{1}{n}\sum_{i<n}x_{iA}\mu_i\right],$$

and analogously $\tilde{\mu}_B$ converges in probability to μ_B. ■

LEMMA 5. *Let $\tilde{\alpha}_A, \tilde{\alpha}_B$ denote the groups random population beliefs. We have that $\tilde{\alpha}_A, \tilde{\alpha}_B$ converge in probability to*

$$\mu_A \hat{\mu}^1 + (1-\mu_A)\hat{\mu}^2,\ \mu_B \hat{\mu}^1 + (1-\mu_B)\hat{\mu}^2,$$
$$\textit{where } \hat{\mu}^1 = \bar{\mu}(\omega_1) + E[\zeta^1],\ \hat{\mu}^2 = \bar{\mu}(\omega_2) + E[\zeta^2].$$

PROOF. Note that

$$\alpha^A = \lim_{n\in\mathbb{N}} \frac{1}{\frac{1}{n}\sum\limits_{i<n} x_{iA}} \frac{1}{n}\sum_{i<n} x_{iA}\left(\mu_i \bar{\mu}(\omega_1) + (1-\mu_i)\bar{\mu}(\omega_2) + \mu_i \zeta_i^1 + (1-\mu_i)\zeta_i^2\right).$$

By Lemma 4, we have that:

$$\frac{1}{\frac{1}{n}\sum\limits_{i<n} x_{iA}} \frac{1}{n}\sum_{i<n} x_i\left(\mu_i \bar{\mu}(\omega_1) + (1-\mu_i)\bar{\mu}(\omega_2)\right) \to_P \mu_A \bar{\mu}(\omega_1) + (1-\mu_A)\bar{\mu}(\omega_2).$$

For the random term

$$\frac{1}{\frac{1}{n}\sum\limits_{i<n} x_{iA}} \frac{1}{n}\sum_{i<n} x_{iA}\left(\mu_i \zeta_i^1 + (1-\mu_i)\zeta_i^2\right)$$

$$\text{let, } L_n = \frac{1}{\frac{1}{n}\sum\limits_{i<n} x_{iA}},\ \text{ and, }\ \hat{K}_n = \frac{1}{n}\sum_{i<n} x_{iA}\left(\mu_i \zeta_i^1 + (1-\mu_i)\zeta_i^0\right).$$

Since the sequences of random variables $\langle x_{iA}\rangle$, $\langle x_{iA}\mu_i\zeta_i^1\rangle$ and $\langle x_{iA}\mu_i\zeta_i^2\rangle$ satisfy the conditions of the Bernstein LLN (conditional on the state ω), it follows that $\langle L_n\rangle$ and $\langle \hat{K}_n\rangle$ converge in probability to L respectively $\hat{K}$ where

$$L = \lim_{i\in\mathbb{N}} E\left[\frac{1}{\frac{1}{n}\sum\limits_{i<n} x_{iA}}\right] \text{ and, } \hat{K} = \lim_{i\in\mathbb{N}} \frac{1}{n}\sum_{i<n} E\left[x_{iA}\mu_i\zeta_i^1\right] + E\left[x_{iA}(1-\mu_i)\zeta_i^2\right].$$

By independence of error terms and first-order beliefs this equals it follows that

$$\hat{K} = \lim_{i\in\mathbb{N}} E[\zeta^1]\left(\frac{1}{n}\sum_{i<n} E\left[x_{iA}\mu_i\right]\right) + E[\zeta^2]\left(\frac{1}{n}\sum_{i<n} E\left[x_{iA}(1-\mu_i)\right]\right).$$

Since $\langle L_n\rangle$ and $\langle \hat{K}_n\rangle$ are uniformly bounded, and by Lemma 4 it follows that the random term

$$\frac{1}{\frac{1}{n}\sum\limits_{i<n} x_{iA}} \frac{1}{n}\sum_{i<n} x_{iA}\left(\mu_i \zeta_i^1 + (1-\mu_i)\zeta_i^2\right)$$

converges in probability, and the claim follows, i.e., $\tilde{\alpha}^A$ converges in probability to

$$\mu_A E[\zeta^1] + (1-\mu_A)E[\zeta^2]$$

concluding the proof. ■

We are now in a position to provide a proof of the main theorem, Theorem 3 (restated below for the reader's convenience).

THEOREM 3. *Suppose the true signal distribution process satisfies Assumptions 1 and 2, and that the agents' knowledge regarding the information structure satisfies Assumption 5. Suppose further that*

$$\left|E[\zeta^1]\right| + \left|E[\zeta^2]\right| < \bar{\mu}\left(\omega_1\right) - \bar{\mu}\left(\omega_2\right). \tag{5}$$

Then, Misspecified Information Population-Mean-Based Aggregation *(Procedure 3) correctly recovers the true state of the world almost surely.*

PROOF. From above, we have that

$$\begin{aligned} \hat{\mu}^1 &= \bar{\mu}(\omega_1) + E[\zeta^1] \\ \hat{\mu}^2 &= \bar{\mu}(\omega_2) + E[\zeta^2]. \end{aligned}$$

We need to show that

$$|\hat{\mu}^{\omega} - \bar{\mu}(\omega)| < |\hat{\mu}^{\neg\omega} - \bar{\mu}(\omega)|$$

for $\omega = \omega_1, \omega_2$. Consider the case of ω_1. Note that by Lemma 5 we have

$$\left|\hat{\mu}^1 - \bar{\mu}(\omega_1)\right| = \left|E[\zeta^1]\right|.$$

For the correct identification of state ω_1 we thus require that

$$\left|\hat{\mu}^2 - \bar{\mu}(\omega_1)\right| > \left|E[\zeta^1]\right|.$$

By Lemma 5 it follows that

$$\left|\hat{\mu}^2 - \bar{\mu}(\omega_1)\right| = \left|\bar{\mu}(\omega_1) - \bar{\mu}(\omega_2) - E[\zeta^2]\right|.(*)$$

Assuming that

$$\left|E[\zeta^2]\right| + \left|E[\zeta^1]\right| < \bar{\mu}(\omega_1) - \bar{\mu}(\omega_2)$$

and plugging into equation $(*)$ yields

$$\left|\hat{\mu}^2 - \bar{\mu}(\omega_1)\right| > \left|\left|E[\zeta^2]\right| + \left|E[\zeta^1]\right| - E[\zeta^2]\right| \geq \left|E[\zeta^1]\right|$$

which establishes correct identification of state ω_1. The argument for the correct identification of state ω_2 is analogous. ■

APPENDIX B. PROOFS FROM SECTION 5

THEOREM 4. *Suppose that $\tilde{\mu}$ has a finite support and the agents report $\tilde{s} = (\tilde{s}_i)_i$. Then there exists a procedure which recovers the "pooled information" posterior on the states, i.e., $\tilde{P}(\cdot|\tilde{s})$.*

PROOF. Assume that $\text{marg}_{\tilde{S}}\tilde{\mu}(\tilde{s}) > 0$. It follows from Theorem III.2.7 of Mertens, Sorin, and Zamir (2015) that the aggregator can derive:

(1) The set $E(\tilde{s})$ which is the smallest set $Y \subseteq \Omega \times \tilde{S}$ of state-hierarchy profiles satisfying
   (a) $(\omega, \tilde{s}) \in Y$ for some $\omega \in \Omega$, and,
   (b) for each $(\omega, \tilde{t}) \in Y$, we have $\{\tilde{t}_i\} \times \text{supp } \tilde{\pi}_i(\tilde{t}_i) \subseteq Y$.

(2) The unique consistent probability $\tilde{\mu}$ on $E(\tilde{s})$.

With (1) and (2), the aggregator can compute $\tilde{\mu}(\cdot|\tilde{s})$. We recap and illustrate both (1) and (2) for the ease of reference.

To construct (1), define $C_i^1(\omega, \tilde{t}) = \{\tilde{t}_i\} \times \text{supp}\tilde{\pi}_i(\tilde{t}_i)$ for each $(\omega, \tilde{t}) \in \Omega \times \tilde{S}$; and inductively, for every $l \geq 1$ and $\tilde{t} \in \tilde{S}$, define

$$C_i^{l+1}(\omega, \tilde{t}) = C_i^l(\omega, \tilde{t}) \bigcup \bigcup_{(\omega', \tilde{t}') \in C_i^l(\omega, \tilde{t})} \bigcup_{j=1}^{N} C_j^1(\omega', \tilde{t}').$$

Then, let

$$C_i(\omega, \tilde{t}) = \bigcup_{l=1}^{\infty} C_i^l(\omega, \tilde{t}).$$

These include $(\omega, \tilde{t})$, the state-hierarchy profiles which agent i regards as possible (i.e., $C_i^1(\omega, \tilde{t})$) at $(\omega, \tilde{t})$, the state-hierarchy profiles which some agent regards at some state-hierarchy profile in $C_i^1(\omega, \tilde{t})$ (i.e., $C_i^2(\omega, \tilde{t})$), and so on.

Consider any ω^* such that $\tilde{\mu}(\omega^*, \tilde{s}) > 0$. Note that $C_i(\omega^*, \tilde{s})$ satisfies properties (a) and (b) above by construction and hence $C_i(\omega^*, \tilde{s}) \supseteq E(\tilde{s})$.[19] Also we can argue that $C_i(\omega^*, \tilde{s}) \subseteq E(\tilde{s})$ inductively.

For (2), it follows from Bayes' rule that for every $(\omega', \tilde{t}') \in C_i^1(\omega, \tilde{t})$ and $\tilde{\mu}(\omega, \tilde{t}) > 0$, we must have $\tilde{t}_i = \tilde{t}_i'$ and hence

$$\frac{\tilde{\mu}(\omega', \tilde{t}')}{\tilde{\mu}(\omega, \tilde{t})} = \frac{\tilde{\pi}_i(\tilde{t}_i')(\omega', \tilde{t}')\,\tilde{\mu}_i(\tilde{t}_i')}{\tilde{\pi}_i(\tilde{t}_i)(\omega, \tilde{t}_{-i})\,\tilde{\mu}_i(\tilde{t}_i)} = \frac{\tilde{\pi}_i(\tilde{t}_i')(\omega', \tilde{t}')}{\tilde{\pi}_i(\tilde{t}_i)(\omega, \tilde{t}_{-i})} > 0.$$

The aggregator knows $\tilde{\pi}_i(\tilde{t}_i)$ and hence can express $\tilde{\mu}(\omega', \tilde{t}')$ as a multiple of $\tilde{\mu}(\omega, \tilde{t})$. Inductively, for every $(\omega', \tilde{t}') \in C_i^l(\omega, \tilde{t})$, $\tilde{\mu}(\omega', \tilde{t}')$ can also be expressed as a multiple of $\tilde{\mu}(\omega, \tilde{t})$.

Hence, $\tilde{\mu}$ is uniquely determined on $E(\tilde{s}) = C_i(\omega^*, \tilde{s})$ since $\tilde{\mu}(\omega^*, \tilde{s}) > 0$. $\tilde{\mu}(\cdot|\tilde{s})$ can thus be calculated. ■

[19]Property (b) is by construction and property (a) follows from the assumption that $\tilde{\mu}(\omega^*, \tilde{s}) > 0$.

B.1. *Eliciting Higher-Order Beliefs up to a Finite Order m*

We adopt his notation here for ease of comparison. Suppose there are two players and two states of nature $\{\omega_1, \omega_2\}$ as in our paper. The construction is easier to explain in terms of the standard partitional model of knowledge.

The model below considers 8 extended states: $\{(\sigma_l, k) : l \in \{1,2\} \text{ and } k \in \{1,2,3,4\}\}$. Interpret state (σ_l, k) in this model as corresponding to ω_l realized for $l = 1, 2$ and any k. In this example, Lipman considers the hierarchy of belief induced by common knowledge that player 1 assigns probability $2/3$ to ω_1 and player 2 assigns probability $1/3$ to ω_1. Such a type does not admit a common prior. However, Lipman shows that the model below, which has a uniform common prior, admits a state $(\sigma_1, 1)$ where the players have the same belief as that of the common knowledge type described above, up to any finite order m.

The prior for the players is the uniform distribution over the extended states, i.e.:

	$(\sigma_1,4)$	$(\sigma_1,3)$	$(\sigma_2,2)$	$(\sigma_1,1)$	$(\sigma_2,1)$	$(\sigma_1,2)$	$(\sigma_2,3)$	$(\sigma_2,4)$
μ	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{1}{8}$

Agents' information is identified with the following partitions:

$$\begin{aligned}
\Pi_1 &= \{\{(\sigma_1,4), (\sigma_1,3), (\sigma_2,2)\} \{(\sigma_1,1), (\sigma_2,1), (\sigma_1,2)\} \ \{(\sigma_2,3), (\sigma_2,4)\}\} \\
\Pi_2 &= \{\{(\sigma_1,4), (\sigma_1,3)\} \{(\sigma_2,2), (\sigma_1,1), (\sigma_2,1)\} \{(\sigma_1,2), (\sigma_2,3), (\sigma_2,4)\}\}.
\end{aligned}$$

Note that each $\Pi_i(\sigma', k)$ is identified with a signal of player i. Observe that player 1 assigns probability one to ω_2 at $\Pi_1(\sigma_2, 4)$ and player 2 assigns probability one to ω_1 at $\Pi_2(\sigma_1, 4)$. Hence, each partition cell of each player induces a different second-order belief. In particular, $(\Pi_1(\sigma_1,1), \Pi_2(\sigma_1,1))$ is the only partition profile at which the second-order beliefs of both players are identical to t. Moreover, conditional on the reported second-order belief at $(\Pi_1(\sigma_1,1), \Pi_2(\sigma_1,1))$ (denoted as $(\Pi_1(\sigma_1,1), \Pi_2(\sigma_1,1))^2$), we have

$$\mu(\omega_1 | (\Pi_1(\sigma_1,1), \Pi_2(\sigma_1,1))) = \frac{1}{2}.$$

Now consider another model modified from the previous one by adding one additional state with the following prior:

	$(\sigma_1,4)'$	$(\sigma_1,3)'$	$(\sigma_2,2)'$	$(\sigma_1,1)'$	$(\sigma_1,1)$	$(\sigma_2,1)$	$(\sigma_1,2)$	$(\sigma_2,3)$	$(\sigma_2,4)$
μ'	$\frac{1}{20}$	$\frac{1}{20}$	$\frac{1}{10}$	$\frac{1}{10}$	0	$\frac{1}{10}$	$\frac{1}{5}$	$\frac{1}{5}$	$\frac{1}{5}$

where $'$ is to indicate that a state is to the left of $(\sigma_1, 1)$ and it will be useful in generalizing the idea to eliciting m orders of beliefs, for $m \geq 3$ in what follows. Agents' information is now given by the partitions:

$$\Pi_1' = \{\{(\sigma_1,4)', (\sigma_1,3)', (\sigma_2,2)', (\sigma_1,1)'\}, \{(\sigma_1,1), (\sigma_2,1), (\sigma_1,2)\} \ \{(\sigma_2,3), (\sigma_2,4)\}\};$$

$$\Pi_2' = \{\{(\sigma_1,4)',(\sigma_1,3)'\},\{(\sigma_2,2)',(\sigma_1,1)',(\sigma_1,1),(\sigma_2,1)\},\{(\sigma_1,2),(\sigma_2,3),(\sigma_2,4)\}\}.$$

Observe now that, conditional on the reported second-order belief at $(\Pi_1'(\sigma_1,1),\Pi_2'(\sigma_1,1))$ (denoted as $(\Pi_1'(\sigma_1,1),\Pi_2'(\sigma_1,1))^2$), we have

$$\mu'\left(\omega_1|\left(\Pi_1'(\sigma_1,1),\Pi_2'(\sigma_1,1)\right)^2\right)=0.$$

Second, observe that

$$\left(\Pi_1'(\sigma_1,1),\Pi_2'(\sigma_1,1)\right)^2=\left(\Pi_1(\sigma_1,1),\Pi_2(\sigma_1,1)\right)^2.$$

This means that the reported second-order beliefs at $(\Pi_1'(\sigma_1,1),\Pi_2'(\sigma_1,1))$ under μ' are the same as at $(\Pi_1(\sigma_1,1),\Pi_2(\sigma_1,1))$ under μ. The basic idea is that in μ', we "shift" the probability which μ assigns to $(\sigma_1,1)$ to the additional state $(\sigma_1,1)'$. This additional state helps us preserve the first-order belief of player 2 at $\Pi_2'(\sigma_1,1)$ at $1/3$, while we decrease the probability $(\sigma_1,4)$ to 0 to preserve the first-order belief of player 1 at $\Pi_1'\left((\sigma_1,4)'\right)$. This takes care of the states to the left of $(\sigma_1,1)$. For states to the right, we double the probability of $(\sigma_2,2)$, $(\sigma_2,3)$, and $(\sigma_2,4)$, so that the first-order belief of player 1 at $\Pi_1'(\sigma_1,1)$ and the first-order belief of player 2 at Π' are also preserved. Therefore, just eliciting the first two orders of beliefs, we cannot distinguish between the model corresponding to μ (under which the posterior would be that both states are equally likely), and μ' (under which the posterior would be that the state is ω_2 for sure).

The construction for $m>2$ is similar but more involved, see below.

B.2. *Construction for $m>2$*

Similarly, we can generalize the construction in Section B.1 to the case with $m\geq 3$ as follows:

$$\Pi_1=\left\{\left\{\{(\sigma_1,2k-1),(\sigma_1,2k),(\sigma_2,k)\}:k=1,...,2^{m-1}\right\},\left\{\left(\sigma_2,2^{m-1}+1\right),....,(\sigma_2,2^m)\right\}\right\};$$
$$\Pi_2=\left\{\left\{\{(\sigma_2,2k-1),(\sigma_2,2k),(\sigma_1,k)\}:k=1,...,2^{m-1}\right\},\left\{\left(\sigma_1,2^{m-1}+1\right),....,(\sigma_1,2^m)\right\}\right\}.$$

$$\mu(\sigma_l,k)=\frac{1}{2^{m+1}},\forall l=1,2,\forall k=1,2,...,2^m.$$

Without loss of generality, assume that $m\geq 3$ is odd. Construct the new model as follows:

$$\Pi_1'=\left\{\begin{array}{c}\{(\sigma_1,1),(\sigma_2,1),(\sigma_1,2)\},\{(\sigma_1,1)',(\sigma_2,2)',(\sigma_1,3)',(\sigma_1,4)'\},\\ \{\{(\sigma_1,2k-1)',(\sigma_1,2k)',(\sigma_2,k)'\}:k=2^{n-1}+1,...,2^n\text{ and }n=3,5,...,m-2\},\\ \{\{(\sigma_1,2k-1),(\sigma_1,2k),(\sigma_2,k)\}:k=2^{n-1}+1,...,2^n\text{ and }n=2,4,...,m-1\},\\ \left\{(\sigma_2,2^{m-1}+1)',....,(\sigma_2,2^m)'\right\}\end{array}\right\};$$

$$\Pi_2' = \left\{ \begin{array}{c} \left\{(\sigma_1,1), (\sigma_2,1), (\sigma_1,1)', (\sigma_2,2)'\right\}, \\ \left\{\left\{(\sigma_2,2k-1)', (\sigma_2,2k)', (\sigma_1,k)'\right\} : k = 2^{n-1}+1, ..., 2^n \text{ and } n = 2,4,...,m-1\right\}, \\ \left\{\left\{(\sigma_2,2k-1), (\sigma_2,2k), (\sigma_1,k)\right\} : k = 2^{n-1}+1, ..., 2^n \text{ and } n = 1,3,...,m-2\right\}, \\ \left\{\left(\sigma_1, 2^{m-1}+1\right), ..., (\sigma_1, 2^m)\right\} \end{array} \right\}.$$

As in the case of $m = 2$, we start from the partition cells $\{(\sigma_1,1), (\sigma_2,1), (\sigma_1,2)\}$ and $\left\{(\sigma_1,1), (\sigma_2,1), (\sigma_1,1)', (\sigma_2,2)'\right\}$ which contain $(\sigma_1,1)$. The states with $'$ are those "to the left" of $(\sigma_1,1)$, whereas the states without $'$ are like those "to the right" of $(\sigma_1,1)$. We can then mimic the idea for $m = 2$ to solve for a prior μ' with the desired properties as follows:

(1) Again, for $x > 0$, set

$$\mu'(\sigma_1,1) = 0 \text{ and } \mu'(\sigma_2,1) = \mu'(\sigma_1,1)' = \mu'(\sigma_2,2)' = x.$$

(2) The number of states with $'$ excluding $(\sigma_1,1)'$ and $(\sigma_2,2)'$ (i.e., states in $\Pi_2(\sigma_1,k)'$ where $k = 2^{n-1}+1, ..., 2^n$ and $n = 2,4,...,m-1$) is :

$$y \equiv 3 \times \left(2^1 + 2^3 + \cdots + 2^{m-2}\right),$$
$$\text{i.e., } y = 2^m - 2.$$

We assign probability $\frac{x}{2}$ to each of these "left" states.

(3) The number of states without $'$ excluding $(\sigma_1,1)$ and $(\sigma_2,1)$ (i.e., $(\sigma_2,2)$ and states in $\Pi_1(\sigma_2,k)$ where $k = 2^{n-1}+1, ..., 2^n$ and $n = 2,4,...,m-1$) is:

$$1 + 3 \times \left(2^1 + 2^3 + \cdots + 2^{m-2}\right) = y + 1.$$

We assign probability $2x$ to each of the "right" states.

(4) Hence for the total probability to sum to 1, we must have:

$$3x + y\frac{x}{2} + (y+1)2x = 1,$$
$$\Longrightarrow x = \frac{2}{10+5y},$$
$$\Longrightarrow x = \frac{1}{5 \times 2^m}.$$

In summary, observe that as desired, we have that

$$\begin{aligned} \mu\left(\omega_1 | \left(\Pi_1(\sigma_1,1), \Pi_2(\sigma_1,1)\right)^m\right) &= \frac{1}{2}. \\ \mu'\left(\omega_1 | \left(\Pi_1'(\sigma_1,1), \Pi_2'(\sigma_1,1)\right)^m\right) &= 0. \\ \left(\Pi_1'(\sigma_1,1), \Pi_2'(\sigma_1,1)\right)^m &= \left(\Pi_1(\sigma_1,1), \Pi_2(\sigma_1,1)\right)^m. \end{aligned}$$

Clearly, we can flip left and right to construct another model Π_1'' and Π_2'' with

$$\begin{aligned} \mu'\left(\omega_1 | \left(\Pi_1''\left(\sigma_1,1\right),\Pi_2''\left(\sigma_1,1\right)\right)^m\right) &= 1; \\ \left(\Pi_1''\left(\sigma_1,1\right),\Pi_2''\left(\sigma_1,1\right)\right)^m &= \left(\Pi_1\left(\sigma_1,1\right),\Pi_2\left(\sigma_1,1\right)\right)^m. \end{aligned}$$

The construction here makes use of the feature that all higher-order beliefs of the common knowledge hierarchy t are degenerate. More precisely, this feature ensures that in "shifting" out the probability which μ assigns to $(\sigma_1, 1)$, we can preserve the higher-order beliefs of $\Pi_1(\sigma_1, 1)$ and $\Pi_2(\sigma_1, 1)$ as long as we can preserve the first-order belief at every partition cell except for $\Pi_1(\sigma_2, 2^m)$ and $\Pi_2(\sigma_1, 2^m)'$.

Appendix C. Counterexamples to SP/ SC for multiple states

Example 1. *Suppose there are three states $\Omega = \{\omega_1, \omega_2, \omega_3\}$, and three signals $S = \{s_1, s_2, s_3\}$. All agents have the initial uniform prior. Agents have a uniform prior over signals and receive conditionally i.i.d. signals, conditional on the state, so that their posteriors can be described as:*

	$P(\omega_1\|\cdot)$	$P(\omega_2\|\cdot)$	$P(\omega_3\|\cdot)$
s_1	0.4	0.21	0.39
s_2	0.45	0.54	0.01
s_3	0.44	0.06	0.5

i.e., the i^{th} row and j^{th} column is the posterior the agent places on ω_j upon signal s_i, $P(\omega_j|s_i)$. Note that this violates the assumption of "diagonal dominance"— $P(\omega_1|s_2) > P(\omega_1|s_1)$, i.e., an agent who sees s_2 (or, indeed s_3) places a higher posterior on ω_1 than an agent who sees signal s_1. However note that an agent who sees signal s_i believes that ω_i is the most likely state.

To be clear, note that this can be achieved by the following distribution of signals conditional on the state (numbers rounded):

	$P(\cdot\|\omega_1)$	$P(\cdot\|\omega_2)$	$P(\cdot\|\omega_3)$
s_1	0.31	0.259	0.433
s_2	0.349	0.667	0.011
s_3	0.341	0.074	0.556

This results, by mechanical calculation, on the following average beliefs in the population:

	$\overline{\mu}(\omega_1)$	$\overline{\mu}(\omega_2)$	$\overline{\mu}(\omega_3)$
ω_1	0.431	0.436	0.422
ω_2	0.274	0.419	0.130
ω_3	0.295	0.145	0.447

here the i^{th} row and and j^{th} column represent the average belief in the population that the state is ω_i when the true state is ω_j.

Note that the resulting expected population average belief is:

	$\alpha(s_1)$	$\alpha(s_2)$	$\alpha(s_3)$
ω_1	0.429	0.434	0.427
ω_2	0.248	0.351	0.211
ω_3	0.323	0.215	0.362

i.e., the column labeled $\alpha(s_i)$ is the expected population average belief of an agent seeing signal s_i.

More generally, we can summarize the set of "surprisingly popular" state(s) in each case as:

	ω_1	ω_2	ω_3
$\alpha(s_1)$	$\omega_1, \underline{\omega_2}$	$\omega_1, \underline{\omega_2}$	ω_3
$\alpha(s_2)$	ω_3	$\omega_1, \underline{\omega_2}$	ω_3
$\alpha(s_3)$	$\omega_1, \underline{\omega_2}$	$\omega_1, \underline{\omega_2}$	ω_3

i.e., the entry in row i corresponding to $\alpha(s_i)$ and column j corresponding to state ω_j is the set of states that are surprisingly popular, in true state ω_j, relative to the expected population average beliefs of an agent who received signal s_i. The underlined entry, if there are multiple, is the one which is surprising by the largest magnitude. For example, relative to an agent seeing s_1 in true state ω_1, both ω_1 and ω_2 are surprisingly popular, but ω_2 is the most surprisingly popular: $(0.274 - 0.248 = 0.028 > 0.02 = 0.431 - 0.429)$.

Note that in our example, the "most surprisingly popular" procedure fails to identify the true state when the state is ω_1, regardless of the signal of who is polled for their expectation about the population average beliefs. Indeed, if an agent who saw signal s_2 is polled, the true state is not even in the set of surprisingly popular states. Note also the failure of first-order stochastic dominance of population average beliefs in this 3-state example—as we argued earlier, that was key to why SP works with binary states, but this can be violated with three or more states.

Another closely related approach proposed is that of "prediction normalized votes" (see Section 1.3 of the Supplementary appendix of Prelec, Seung, and McCoy (2017)): each agent votes for the state they believe is more likely (which, in this example, was the signal they saw, i.e., an agent votes for state ω_i if they see signal s_i). The fraction of votes each state ω_j receives is normalized by the sum, over all states ω_k, the ratio of predicted vote fraction for ω_k by a ω_j voter to the predicted vote fraction for ω_j by an ω_k voter. A simple calculation shows that given the specific numerical assumption we constructed, ω_2 will have a higher prediction normalized vote than ω_1 when the true state is ω_1.