A Population's Feasible Posterior Beliefs*

Itai Arieli†, Yakov Babichenko‡

May 22, 2022

Abstract

We consider a population of Bayesian agents who share a common prior over some finite state space and each agent is exposed to some information about the state. We ask which distributions over empirical distributions of posteriors beliefs in the population are feasible. We provide a necessary and sufficient condition for feasibility. We apply this result in several domains. First, we study the problem of maximizing the polarization of beliefs in a population. Second, we provide a characterization of the feasible agent-symmetric product distributions of posteriors. Finally, we study an instance of a private Bayesian persuasion problem and provide a clean formula for the sender's optimal value.

1 Introduction

Economics is deeply concerned with the question of whether agents' behaviors can be rationalized in the sense that they can be explained as an equilibrium behavior in a Bayesian model with rational agents. One such related question is under what conditions agents' beliefs, i.e., posterior distribution, are compatible with Bayes's rule? This question is well understood for the single agent where by the splitting lemma ([Aumann

*Itai acknowledges support from the Ministry of Science and Technology (grant 19400214). Yakov acknowledges support from the Israel Science Foundation (grant 2021296).

†Faculty of Industrial Engineering and Management, Technion: The Israel Institute of Technology, iarieli@technion.ac.il

‡Faculty of Industrial Engineering and Management, Technion: The Israel Institute of Technology, yakovbab@technion.ac.il

et al., 1995, Blackwell, 1951]) feasibility simply means that the expectation of the posterior equals the prior. But the characterization of feasible posterior distribution for two or more agents is trickier.

The first paper to address the feasibility of posterior distributions for more than one agent is [Dawid et al., 1995] who characterize feasible distributions for two agents. Recently a growing number of papers have examined the feasibility of posterior distributions and their connection to information design problems. [Arieli et al., 2021] connect the notion of feasibility to the agreement theorem of [Aumann, 1976] and to the related no-trade theorem of [Milgrom and Stokey, 1982]. A follow-up paper by [Morris, 2020] generalizes the result in [Arieli et al., 2021] to an arbitrary finite state space. Despite these elegant connections, in general, determining whether a given distribution of posterior beliefs is feasible is a tough question.

In this work we study an *anonymous* variant of the feasibility question. The object whose feasibility we are willing to determine is a distribution over the *empirical distributions of posterior beliefs in the population* rather then a distribution over *profiles of posteriors*. The difference between these two variants is that a profile of posteriors indicates which agent holds which posterior, while the anonymous version of the problem only reveals how many agents hold each posterior without revealing their identity.

Surprisingly, we show that determining feasibility for the anonymous variant of the problem is relatively easy, especially in the case where the number of different posteriors for an individual in the population is low. Our characterization (Theorem 1) provides a necessary and sufficient condition for a distribution over empirical samples of a population to be feasible. The characterization relies on the observation that conditional on a state the frequencies of posterior beliefs in the population can be identified. We show that every combination of spreads of these expected frequencies (one for each state) is feasible. Unlike the feasibility characterization by [Arieli et al., 2021] and [Morris, 2020], our characterization does not become more complex as the number of agents in the population increases. Furthermore, the complexity of our characterization does not increase with the number of states either. However, the parameter of the number of possible posteriors in the population has an effect on the complexity of our characterization.

As was mentioned above, the feasibility problem is an intriguing instance of rationalization in the case where an econometrician observes the distribution over empirical

distributions of posteriors in the population. In the context of rationalization, it is natural to consider another variant where the econometrician observes a single sample; i.e., she observes a single empirical distribution of posteriors and wants to determine whether it is possible that the agents are Bayesian.[1] [Molavi, 2021] shows that every empirical distribution of posteriors can be rationalized. This question is equivalent to understanding which empirical distributions of posteriors might belong to the support of a feasible distribution. Our results can be viewed as an extension of [Molavi, 2021]; however, we go a step further and determine the entire set of feasible distributions rather than focus on points that might belong to the support.

We apply our main theorem in three different domains.

Polarization The question of how polarized the posteriors of two Bayesian agents might be has been previously studied. Different measures of polarization have been suggested, including the distance between the posteriors [Dubins and Pitman, 1980], the distance to the power α [Arieli et al., 2021], and the indicator function of the posterior's distance being above some threshold [Burdzy and Pal, 2019, Burdzy and Pitman, 2020]. The case of a quadratic distance can be naturally generalized to multiple agents by considering the variance of the belief's in the population. The maximal polarization for two agents is obtained by revealing the state to one agent and providing no information to the other agent ([Arieli et al., 2021]). We extend this result to every population of agents of even size: the maximal polarization is obtained by revealing the state to half of the population and providing no information to the other half. Furthermore, we extend this result to an arbitrary state space (not necessarily binary).

Private private information A fascinating instantiation of the feasibility problem is that of product distributions. If the distribution of belief profiles is a product distribution then the information that agent i receives does not affect her belief about the posteriors of her peers. [He et al., 2021] refer to this property as *private private information* of the agents. The feasibility of product distributions has been considered in [He et al., 2021, Arieli et al., 2021, Gutmann et al., 1991a] and we currently have a good understanding of this problem for the two-agent case. However, the feasibility of product distributions for

[1] [Shmaya and Yariv, 2016, De Oliveira and Lamba, 2019] have studied Bayesian rationalization of the behavior of a single agent in a repeated environment.

more then two agents currently does not have a characterization. We contribute to this line of literature by providing a characterization of feasible *agent-symmetric* product distributions for an arbitrary number of agents. This characterization takes a simple form in the case where the marginals of the product distributions are binary-supported.

Private Bayesian Persuasion The Bayesian persuasion setting is quite well understood in the case where the sender communicates with the receivers publicly. In particular the value of the sender is characterized by the notion of concavification [Kamenica and Gentzkow, 2011, Aumann et al., 1995]. However, if the sender is allowed to communicate privately with each receiver, the characterization of the value is known only in special cases. A particular case that has been studied in [Arieli and Babichenko, 2019, Babichenko and Barman, 2016, Dughmi and Xu, 2019] is the following. A marketer (sender) tries to persuade agents (receivers) to adopt a product whose quality is either high or low. [Arieli and Babichenko, 2019, Babichenko and Barman, 2016, Dughmi and Xu, 2019] characterize optimal policies for the sender in special cases of sender's utility function such as supermodular or anonymous submodular utilities. We provide an additional natural instance of the problem where the optimal sender's value (and the optimal policy) can be characterized. We consider the case with a homogeneous population of agents (i.e., all agents have an identical utility function) and the sender's utility function is a monotonic anonymous function (i.e., a function of a fraction of the adopters). Interestingly, similarly to the standard Bayesian persuasion model ([Kamenica and Gentzkow, 2011]), our characterization also relies on the notion of concavification. However, the concavification in our case is applied in a different domain.

We illustrate the key ideas of our main result via an example in Section 2. Section 3 proceeds with the model and the main result. Section 4 studies the implications of our results to maximal polarization of beliefs in a population. In Section 5 we discuss the connections of our anonymous version of feasibility with the previously studied (non-anonymous) version and present an application to private private information. Finally, Section 6 demonstrates an application in a Bayesian persuasion framework.

2 Warm-up Example

Let $\Omega = \{0, 1\}$ be a binary state space with a common prior $\mu = \frac{1}{2}$. We identify $\Delta(\Omega)$ with $[0, 1]$, which indicates the probability that is assigned to $\omega = 1$. An information structure for n agents comprises n measurable sets $\{S_i\}_{i=1}^n$ as well as a probability distribution $G \in \Delta(\Omega \times S_1 \times \cdots \times S_n)$. We note that almost surely any realization of n signals $(s_1, \ldots, s_n)$ defines n posterior beliefs $G(\omega|s_i)$ for $i \in [n]$. In the example of the present section we assume that $G(\omega|s_i) \in \{\frac{1}{4}, \frac{3}{4}\}$ almost surely for every $i \in [n]$; namely, any information that an agent is exposed to results in a posterior of either $\frac{1}{4}$ or $\frac{3}{4}$. The information structure G and the signal realization $s = (s_1, \ldots, s_n)$ induce a *distribution of posteriors in the population*. Since we have assumed that there are only two possible posteriors we can identify the distribution of posteriors in the population with a fraction $H_s \in \{0, \frac{1}{n}, \frac{2}{n}, ..., 1\}$, where $H_s = \frac{k}{n}$ indicates that k agents in the population have the high posterior $\frac{3}{4}$, and $n - k$ agents have the low posterior $\frac{1}{4}$.

The information structure G induces a distribution over distributions of posteriors in the population $\mathbf{P}_G \in \Delta(\{0, \frac{1}{n}, \frac{2}{n}, ..., 1\})$. We will call such a $\mathbf{P}_G$ the *population's empirical posteriors* for short. The population's empirical posteriors anonymize the identity of the individuals and considers only the empirical distributions of the belief distribution.

A population's empirical posteriors $\mathbf{P}$ is called *feasible* if $\mathbf{P} = \mathbf{P}_G$ for some information structure of n agents. A characterization of the feasible population's empirical posteriors is given in the following proposition. We denote by δ_c the Dirac measure on an element c.

Proposition 1. A population's empirical posteriors $\mathbf{P} \in \Delta(\{0, \frac{1}{n}, \frac{2}{n}, ..., 1\})$ is feasible if and only if $\mathbf{P}$ is a mean-preserving spread of the distribution $\frac{1}{2}\delta_{\frac{1}{4}} + \frac{1}{2}\delta_{\frac{3}{4}}$.

To demonstrate the applicability of Proposition 1 consider, for instance, a population of $n = 9$ agents. The uniform distribution $U(\{0, \frac{1}{9}, \frac{2}{9}, ..., 1\})$ is feasible because $U(\{0, \frac{1}{9}, \frac{2}{9}, ..., 1\})$ is a mean-preserving spread of $\frac{1}{2}\delta_{\frac{1}{4}} + \frac{1}{2}\delta_{\frac{3}{4}} = U(\{\frac{1}{4}, \frac{3}{4}\})$; that is, there exists an information structure such that with probability $\frac{1}{10}$ the distribution of posteriors in the population will be i agents with posterior $\frac{3}{4}$ and $9 - i$ agents with posterior $\frac{1}{4}$ for every $i = 0, 1, ..., 9$.

The uniform distribution $U(\{\frac{1}{9}, \frac{2}{9}, ..., \frac{8}{9}\})$ is infeasible because $U(\{\frac{1}{9}, \frac{2}{9}, ..., \frac{8}{9}\})$ is not a mean-preserving spread of $\frac{1}{2}\delta_{\frac{1}{4}} + \frac{1}{2}\delta_{\frac{3}{4}} = U(\{\frac{1}{4}, \frac{3}{4}\})$; that is, there is no information

structure such that with probability $\frac{1}{8}$ the distribution of posteriors in the population will be i agents with posterior $\frac{3}{4}$ and $9-i$ agents with posterior $\frac{1}{4}$ for every $i = 1, 2, ..., 8$.

Proposition 1 is a special case of our main result, which characterizes the feasibility of an arbitrary distribution. We find it useful to present the proof of Proposition 1 because it contains the main proof ideas of the general case (Theorem 1).

Proof of Proposition 1. First, we show that every feasible distribution must be a mean-preserving spread of $\frac{1}{2}\delta_{\frac{1}{4}} + \frac{1}{2}\delta_{\frac{3}{4}}$. Conditional on state $\omega = 1$, every agent $i \in [n]$ has a posterior of $\frac{3}{4}$ with probability $\frac{3}{4}$. Therefore, conditional on state $\omega = 1$, the expected number of agents with the high posterior $\frac{3}{4}$ is $\frac{3}{4}$. Similarly, conditional on state $\omega = 0$, the expected number of agents with the high posterior $\frac{3}{4}$ is $\frac{1}{4}$. Denoting by $\mathbf{P}^\omega$ the population's posteriors conditional on ω we obtain a decomposition of $\mathbf{P}$ as $\mathbf{P} = \frac{1}{2}\mathbf{P}^0 + \frac{1}{2}\mathbf{P}^1$ where $\mathbb{E}[\mathbf{P}^0] = \frac{1}{4}$ and $\mathbb{E}[\mathbf{P}^1] = \frac{3}{4}$. We refer to $\mathbf{P}^0$ as the spread of $\delta\frac{1}{4}$, we refer to $\mathbf{P}^1$ as the spread of $\delta\frac{3}{4}$, and we get that $\mathbf{P}$ is a mean-preserving spread of $\frac{1}{2}\delta_{\frac{1}{4}} + \frac{1}{2}\delta_{\frac{3}{4}}$.

Conversely, let $\mathbf{P} = \frac{1}{2}\mathbf{Q}^0 + \frac{1}{2}\mathbf{Q}^1$ where $\mathbb{E}[\mathbf{Q}^0] = \frac{1}{4}$ and $\mathbb{E}[\mathbf{Q}^1] = \frac{3}{4}$ be a mean-preserving spread of $\frac{1}{2}\delta_{\frac{1}{4}} + \frac{1}{2}\delta_{\frac{3}{4}}$. Since $\mathbf{P} \in \Delta(\{0, \frac{1}{n}, ..., 1\})$ we also have $\mathbf{Q}^\omega \in \Delta(\{0, \frac{1}{n}, ..., 1\})$. Denote $\alpha_k^\omega = \mathbf{Q}^\omega(\{\frac{k}{n}\})$. We define an information structure G as follows. In state ω we draw an integer k with probability α_k^ω, then we choose a subset $A \subset [n]$ of size $|A| = k$ uniformly at random (among all possible subsets of size k), and finally we send a high signal h to all agents in the subset A and a low signal l to the remaining agents. We need to show that the posterior of an agent who has observed a signal h is $\frac{3}{4}$ and the posterior of an agent who has observed a signal l is $\frac{1}{4}$. In state ω the signal h is obtained with probability $\sum_{i=0}^{n} \alpha_k^\omega \frac{k}{n} = \mathbb{E}[\mathbf{Q}^\omega]$ because α_k^ω is the probability that the subset is of size k and since the subset is random the probability that agent i is included is $\frac{k}{n}$. Therefore, the posterior of an agent who has observed signal h equals

$$\frac{\mathbb{E}[\mathbf{Q}^1]}{\mathbb{E}[\mathbf{Q}^0] + \mathbb{E}[\mathbf{Q}^1]} = \frac{3}{4}.$$

Similar considerations show that the posterior of an agent who has observed a signal l is $\frac{1}{4}$. □

3 Model and the Main Result

It turns out that the characterization of feasibility from Section 2 can be naturally extended to a much more general case. Neither the restriction to binary possible posteriors nor the restriction to a binary state space are needed. We start by introducing the general model.

For a set A we denote by $\Delta_n(A) \subset \Delta(A)$ the set of empirical samples of size n, i.e., $\Delta_n(A) = \{\sum_{k\in[n]} \frac{1}{n}\delta_{a^k} : a^k \in A \text{ for } k \in [n]\}$ (note that any $a \in A$ may appear in the sum more than once).

Let $\Omega = \{1, \ldots, m\}$ be a finite state space. An information structure for n agents comprises n measurable sets $\{S_j\}_{j=1}^n$ as well as a probability distribution $G \in \Delta(\Omega \times S_1 \times \cdots \times S_n)$. We denote by $\mathcal{G}_n$ the set of all information structures of n agents. The marginal distribution on Ω forms the prior; i.e., the prior $\mu \in \Delta(\Omega)$ satisfies $\mu_i = G(\omega = i)$ for any state $\omega \in \Omega$. Any realization of n signals $(s_1, \ldots, s_n)$ defines almost surely n posterior beliefs $G(\omega|s_j)$ for $j \in [n]$. Hence G and the signal realization $s = (s_1, \ldots, s_n)$ induce a *distribution of posteriors in the population* $H_s \in \Delta_n(\Delta(\Omega))$ that is given by $\sum_{j\in[n]} \frac{1}{n}\delta_{G(\omega|s_j)}$.

The information structure G induces a distribution over distributions of posteriors in the population $\mathbf{P}_G \in \Delta(\Delta_n(\Delta(\Omega)))$, which we call the *population's empirical posteriors* for short.

In the case where S_j is finite for every j, we let $G(s)$ be the unconditional probability of a signal vector $s \in S_1 \times \cdots \times S_n$. In this case we can write

$$\mathbf{P}_G = \sum_{s\in S_1\times\cdots\times S_n} G(s)\delta_{H_s}, \tag{1}$$

where $\delta_{H_s} \in \Delta_n(\Delta(\Omega))$ is a Dirac measure on H_s.

Definition 1. A distribution $\mathbf{P} \in \Delta(\Delta_n(\Delta(\Omega)))$ is called *feasible* for a prior $\mu \in \Delta(\Omega)$ if $\mathbf{P} = \mathbf{P}_G$ for some information structure $G \in \mathcal{G}_n$. We denote by $\mathcal{P}_n(\mu) \subseteq \Delta(\Delta_n(\Delta(\Omega)))$ the set of all feasible distributions for a prior $\mu \in \Delta(\Omega)$ and n agents.

Note that if the signal sets $\{S_j\}_{j=1}^n$ are finite, $\mathbf{P}_G$ assigns probability one to measures $H_s \in \Delta_n(\Delta(\Omega))$ that are supported on some finite set $\{y_1, \ldots, y_m\} \subseteq \Delta(\Omega)$ (where $m \leq \sum_j |S_j|$). In general, we say that $\mathbf{P} \in \mathcal{P}$ is *finitely supported* if there exists a finite

set $Y \subseteq \Delta(\Omega)$ of posteriors such that $\mathbf{P}$ assigns probability one to the set of measures $\Delta(\Delta_n(Y))$.

3.1 Main Result

For any $\mathbf{P} \in \Delta(\Delta_n(\Delta(\Omega)))$ we let $\tilde{\mathbf{P}} \in \Delta(\Delta(\Omega))$ be the *expected* measure that $\mathbf{P}$ induces on $\Delta(\Omega)$. That is, for every Borel subset $B \subseteq \Delta(\Omega)$,

$$\tilde{\mathbf{P}}(B) = \int_{\Delta(\Omega)} H(B) \mathrm{d}\mathbf{P}(H).$$

In particular, if $\mathbf{P}$ is finitely supported we can write $\mathbf{P} = \sum_{j=1}^{k} a_j H_j$ (where $H_j \in \Delta_n(\Delta(\Omega))$, $a_j \geq 0$ and $\sum_j a_j = 1$); then

$$\tilde{\mathbf{P}}(B) = \sum_{j=1}^{k} a_j H_j(B).$$

We define m other distributions $\tilde{\mathbf{P}}^{\omega} \in \Delta(\Delta(\Omega))$ for $\omega \in \Omega$ as follows:

$$\mathrm{d}\tilde{\mathbf{P}}^{\omega}(x) = \frac{1}{\mu_\omega} x_\omega \mathrm{d}\tilde{\mathbf{P}}(x).$$

If $\mathbf{P}$ is finitely supported we simply have $\tilde{\mathbf{P}}^{\omega}(x) = \frac{x_\omega \tilde{\mathbf{P}}(x)}{\mu_\omega}$.

Finally, let $\mathbf{P}' \in \Delta((\Delta(\Delta(\Omega)))$ be the the measure that is defined as follows:

$$\mathbf{P}' = \sum_{\omega \in [m]} \mu_\omega \delta_{\tilde{\mathbf{P}}^{\omega}},$$

where $\delta_{\tilde{\mathbf{P}}^{\omega}} \in \Delta((\Delta(\Delta(\Omega)))$ are Dirac measures. Our main result states the following:

Theorem 1. A distribution $\mathbf{P} \in \Delta(\Delta_n(\Delta(\Omega)))$ is feasible for n agents and a prior $\mu \in \Delta(\Omega)$ if and only if $\mathbf{P}$ is a mean-preserving spread of $\mathbf{P}'$.

The proof follows a similar logic to the proof of Proposition 1; however, the general case differs from the binary-posterior case in the following respect. A key observation in the proof of Theorem 1, as well as the proof of Proposition 1, is that the expected distribution of posteriors conditional on a state ω can be identified. In the binary-posterior case, conditional on any state ω, *every individual* in the population has on

average the high posterior ($\frac{3}{4}$ in Proposition 1) with a fixed probability that is determined by the two possible posteriors. In the case of more than two posteriors we cannot deduce such a strong statement about each individual in the population.[2] Nevertheless, it turns out that the *average* distribution of posteriors over the population conditional on a state can be identified. To do so, we simply do the exercise of deducing the average *conditional* ($\tilde{\mathbf{P}}^{\omega}$) distribution of posteriors from the average *ex-ante* distribution of posteriors ($\tilde{\mathbf{P}}$). The validity of the above deduction follows from the linearity of expectation and the linearity of the operator that maps an ex-ante distribution of posteriors to a conditional distribution of posteriors. Once we show that conditional on a state $\omega \in \Omega$ the average distribution of posteriors must be $\tilde{\mathbf{P}}^{\omega}$, the necessity for the mean-preserving spread (of $\mathbf{P}'$) condition immediately follows. To prove sufficiency we construct an information structure with the population's empirical posteriors $\mathbf{P}$ using similar arguments to those of Proposition 1.

3.2 The binary-state binary-posterior case

In the case where the state space is binary and the posteriors get only two values a, b where $0 \leq a < \mu < b \leq 1$, the distribution $\mathbf{P}$ is an element of $\Delta(\Delta_n(\{a, b\}))$ or simply an element of $\Delta(\{0, \frac{1}{n}, ..., 1\})$ that indicates the weight of b in the distribution. In the latter representation of the population's empirical posteriors the terms $\tilde{\mathbf{P}}^{\omega}$ for $\omega = 0, 1$ represent the expected fraction of agents that have the high posterior conditional on a state. This expectation equals the probability of an agent having the high posterior in the *single-agent* case. Simple calculations show that $\tilde{\mathbf{P}}_1 = \frac{(\mu-a)b}{(b-a)\mu}$ and $\tilde{\mathbf{P}}_0 = \frac{(\mu-a)(1-b)}{(b-a)(1-\mu)}$.

Thus, we deduce the following corollary:

Corollary 1. A population's empirical posteriors $\mathbf{P} \in \Delta(\{0, \frac{1}{n}, \frac{2}{n}, ..., 1\})$ is feasible for the prior μ and the posteriors $a < \mu < b$ if and only if $\mathbf{P}$ is a mean-preserving spread of

$$\mu\delta_{\frac{(\mu-a)b}{(b-a)\mu}} + (1-\mu)\delta_{\frac{(\mu-a)(1-b)}{(b-a)(1-\mu)}}.$$

The fact that μ_P and $\tilde{\mathbf{P}}^{\omega}$ can be deduced directly from the values of the posteriors a

[2] Consider, for instance, a binary-state space with three possible posteriors $0 < a < b < \mu < c < 0$. Agent 1 has only two possible posteriors a and c. Agent 2 has only two possible posteriors b and c. Obviously, agents 1 and 2 will have different distributions of posteriors conditional on any state.

and b is a special feature of the binary-posterior case.[3] In the general case, the terms μ_P and $\tilde{\mathbf{P}}^\omega$ are not determined by the support of the posteriors but rather are determined by the distribution $\mathbf{P}$ as described above in the definitions of these terms.

We also notice that for single-dimensional distributions it is easy to determine whether the distribution is a mean-preserving spread of a given binary-support distribution as follows. Let $\mathbf{P} \in \Delta(\{0, \frac{1}{n}, \frac{2}{n}, ..., 1\})$ be a distribution that assigns a probability of p_i to $\frac{i}{n}$ and let $\alpha\delta_a + (1-\alpha)\delta_b$ be a binary-support distribution with $a < b$. We consider the *α-quantile distribution of* $\mathbf{P}$ that is defined as follows. We let k be the index such that $\sum_{i<k} p_i \leq \alpha < \sum_{i\leq k} p_i$. We let $q_k = \alpha - \sum_{i<k} p_i$. The α-quantile distribution of $\mathbf{P}$ is denoted by $\mathbf{P}|_\alpha$ and is defined to be the distribution that assigns a probability of $\frac{p_i}{\alpha}$ to $\frac{i}{n}$ for every $i < k$ and assigns a probability of $\frac{q_k}{\alpha}$ to $\frac{k}{n}$. Simply speaking, $\mathbf{P}|_\alpha$ is the distribution of $\mathbf{P}$ conditional on the event that the realization belongs to the lower α-quantile of $\mathbf{P}$. The following lemma characterizes the mean-preserving spread property.

Lemma 1. A distribution $\mathbf{P} \in \Delta(\{0, \frac{1}{n}, \frac{2}{n}, ..., 1\})$ is a mean-preserving spread of $\mathbf{B} = \alpha\delta_a + (1-\alpha)\delta_b$ if and only if $\mathbb{E}[\mathbf{P}] = \mathbb{E}[\mathbf{B}]$ and $\mathbb{E}[\mathbf{P}|_\alpha] \leq a$.

Proof. We denote by $F_{\mathbf{P}}, F_{\mathbf{B}}$ the CDFs for $\mathbf{P}$ and $\mathbf{B}$. We have that $\mathbf{P}$ is a mean preserving spread of $\mathbf{B}$ if and only if $\mathbb{E}[\mathbf{P}] = \mathbb{E}[\mathbf{B}]$ and for every $t \in [0,1]$ we have $\int_0^t F_{\mathbf{P}}(x)dx \geq \int_0^t F_{\mathbf{B}}(x)dx$. Moreover, it is well known that it suffices to check this inequality only for points t such that $F_{\mathbf{P}}(t)$ crosses $F_{\mathbf{B}}(t)$. In our case $F_{\mathbf{B}}$ is piecewise constant with values $0, \alpha, 1$ in the corresponding segments $[0,a), [a,b), [b,1]$. Therefore, in our case it is sufficient to check the inequality at the single point $t = \frac{k}{n}$.

Note that the CDF of $\mathbf{P}|_\alpha$ is $F_{\mathbf{P}|_\alpha}(x) = \alpha^{-1}F_{\mathbf{P}}(x)$. Hence we have

$$\int_0^{\frac{k}{n}} F_{\mathbf{P}}(x)dx = \alpha \int_0^{\frac{k}{n}} F_{\mathbf{P}|_\alpha}(x)dx = \alpha\left(\frac{k}{n} - \int_0^{\frac{k}{n}} (1 - F_{\mathbf{P}|_\alpha}(x))dx\right) = \alpha\left(\frac{k}{n} - \mathbb{E}[\mathbf{P}|_\alpha]\right).$$

We also have $\int_0^{\frac{k}{n}} F_{\mathbf{B}}(x)dx = \alpha(\frac{k}{n} - a)$ and therefore $\int_0^{\frac{k}{n}} F_{\mathbf{P}}(x)dx \geq \int_0^{\frac{k}{n}} F_{\mathbf{B}}(x)dx$ if and only if $\mathbb{E}[\mathbf{P}|_\alpha] \leq a$. □

[3] Those properties hold for the case of binary posteriors and an arbitrary state space.

3.3 Proof of Theorem 1

Let $\mathbf{P} \in \Delta(\Delta_n(\Delta(\Omega))$ be a mean-preserving spread of $\mathbf{P}'$. Namely, there exist m distributions $\mathbf{Q}^\omega \in \Delta(\Delta(\Omega))$ with expectation $\tilde{\mathbf{P}}^\omega$ such that

$$\mathbf{P} = \sum_{\omega \in \Omega} \mu_\omega \mathbf{Q}^\omega.$$

We now define a signaling structure G that implements $\mathbf{P}$ as follows. Let $S_j = \Delta(\Omega)$ for every agent j and let $G(\omega) = \mu_i$ for every state $i \in \Omega$. We define $G(s|\omega)$ by coupling this measure with $\mathbf{Q}^\omega$ as follows. We draw $H \in \Delta_n(\Delta(\Omega))$ according to $\mathbf{Q}^\omega$. Every such H has the form $H = \sum_{k \in [n]} \frac{1}{n} \delta_{x^k}$ where $x^k \in \Delta(\Omega)$. For every realized H we draw a permutation $\pi : [n] \to [n]$ uniformly at random across all permutations and send signal $x^{\pi(j)}$ to every agent j. Thus, overall the probability that agent j observes the signal x^k is $\frac{1}{n}$ for every $j, k \in [n]$.

We claim that $\mathbf{P}_G = \mathbf{P}$. To see this, by the construction, it is enough to show that $G(\omega|s_j = x) = x_\omega$ for every agent $j \in [n]$ and state $\omega \in \Omega$. We note first that since H is drawn according to $\mathbf{Q}^\omega$ and since the expectation of $\mathbf{Q}^\omega$ is $\tilde{\mathbf{P}}^\omega$ we have that $G(s_j \in B|\omega) = \tilde{\mathbf{P}}^\omega(B)$ for every Borel measurable set $B \subset \Delta(\Omega)$. Therefore, by Bayes's rule, for every $\omega \in \Omega$, we have that

$$G(\omega|s_j = x) = \frac{\mu_\omega \frac{\mathrm{d}\tilde{\mathbf{P}}^\omega}{\mathrm{d}\tilde{\mathbf{P}}}(x)}{\sum_{\omega' \in \Omega} \mu_{\omega'} \frac{\mathrm{d}\tilde{\mathbf{P}}^{\omega'}}{\mathrm{d}\tilde{\mathbf{P}}}(x)} \tag{2}$$

with probability 1.

We recall that $\mathrm{d}\tilde{\mathbf{P}}^\omega(x) = \frac{x_\omega}{\mu_\omega} \mathrm{d}\tilde{\mathbf{P}}(x)$. Therefore, $\frac{\mu_{\omega'} \mathrm{d}\tilde{\mathbf{P}}_{\omega'}(x)}{\mathrm{d}\tilde{\mathbf{P}}(x)} = x_{\omega'}$ for every state $\omega' \in \Omega$. Hence, Equation (2) implies that $G(\omega|s_j = x) = \frac{x_\omega}{\sum_{\omega' \in \Omega} x_{\omega'}} = x_\omega$ as desired.

We now turn to the converse direction; namely, we show that every feasible $\mathbf{P} \in \mathcal{P}_n$ is a mean-preserving spread of $\mathbf{P}'$. For clarity of exposition we prove this statement only for the case where $\mathbf{P}$ is finitely supported. The extension of the proof to arbitrary distributions $\mathbf{P}$ is relegated to the Appendix A and essentially follows from the fact that any distribution $\mathbf{P}$ can be approximated (in the weak* topology) by finitely supported ones.

We assume that $\mathbf{P} = \mathbf{P}_G$ is induced by an information structure $G \in \Delta(\Omega \times S_1 \times \cdots \times S_n)$ with finite signal sets S_i for every $i \in [n]$ (the finite support assumption). In addition

to the distribution $\mathbf{P}$, we next define a few other auxiliary measures. First, for every state $\omega \in \Omega$ let $\mathbf{P}_G^\omega \in \Delta(\Delta_n(\Delta(\Omega)))$ be the conditional distribution of H_s conditional on state ω. That is, similar to Equation (1),

$$\mathbf{P}_G^\omega = \sum_{s \in S_1 \times \cdots \times S_n} G(s|\omega)\delta_{H_s}.$$

It follows that

$$\mathbf{P} = \sum_{\omega \in \Omega} \mu_\omega \mathbf{P}_G^\omega. \tag{3}$$

For every $\omega \in \Omega$, let $\tilde{\mathbf{P}}_G^\omega$ be the expectation of $\mathbf{P}_G^\omega$. By definition, the measure $\mathbf{P}_G^\omega$ is a mean-preserving spread of $\delta_{\tilde{\mathbf{P}}_G^\omega}$. Therefore, in particular, by Equation (3) $\mathbf{P}$ is a mean-preserving spread of the measure $\sum_{\omega \in \Omega} \mu_\omega \delta_{\tilde{\mathbf{P}}_G^\omega}$. In order to complete the proof we need to show that $\tilde{\mathbf{P}}_\omega$ as defined above equals $\tilde{\mathbf{P}}_G^\omega$ for every state $\omega \in \Omega$.

For each signal $s_j \in S_j$ let $G(\ |s_j) \in \Delta(\Omega)$ be the conditional probability on Ω given the signal s_j. Let $H_{s_j} = \delta_{G(\ |s_j)}$. Using this notation, for any vector of realized signals $s = (s_1, \ldots, s_n)$, one can write the measure H_s as follows:

$$H_s = \frac{1}{n} \sum_{j=1}^{n} H_{s_j}.$$

Therefore, we have

$$\mathbf{P} = \sum_{s \in S_1 \times \cdots \times S_n} G(s) \frac{1}{n} \sum_{j=1}^{n} H_{s_j}.$$

Hence, by definition of $\tilde{\mathbf{P}}$, for every $x \in \Delta(\Omega)$ we have

$$\tilde{\mathbf{P}}(\{x\}) = \sum_{s \in S_1 \times \cdots \times S_n} G(s) \frac{1}{n} \sum_{j=1}^{n} H_{s_j}(x) = \frac{1}{n} \sum_{j=1}^{n} \sum_{s_j \in S_j, G(\ |s_j)=x} G(s_j).$$

Similarly, for any state $\omega \in \Omega$,

$$\tilde{\mathbf{P}}_G^\omega(\{x\}) = \sum_{s \in S_1 \times \cdots \times S_n} G(s|\omega) \frac{1}{n} \sum_{j=1}^{n} H_{s_j}(x) = \frac{1}{n} \sum_{j=1}^{n} \sum_{s_j \in S_j, G(\ |s_j)=x} G(s_j|\omega).$$

It follows from Bayes's rule that if $G(\omega|s_j) = x_\omega$, then $G(s_j|\omega) = x_\omega \frac{G(s_j)}{\mu_\omega}$. Hence $\tilde{\mathbf{P}}_G^\omega(\{x\}) = \frac{x_\omega}{\mu_\omega} \tilde{\mathbf{P}}(\{x\})$ for every $x \in \Delta(\Omega)$. Thus $d\tilde{\mathbf{P}}_{G,\omega}(\{x\}) = \frac{x_\omega}{\mu_\omega} d\tilde{\mathbf{P}}(\{x\})$. Therefore, for every $\omega \in \Omega$ it holds that $\tilde{\mathbf{P}}_G^\omega = \tilde{\mathbf{P}}^\omega$, as desired.

4 Polarization

The question of how polarized the posteriors of Bayesian agents might be has been studied in [Dubins and Pitman, 1980], [Burdzy and Pitman, 2020], and [Arieli et al., 2021]. These papers focuse on the case of two receivers. In [Dubins and Pitman, 1980] polarization is measured by the absolute distance $|x_1 - x_2|$ and more generally in [Arieli et al., 2021] polarization is measured by the function $|x_1 - x_2|^\alpha$ for some $\alpha > 0$, where $x_i \in [0, 1]$ denotes the posteriors of agent $i = 1, 2$. The case of $\alpha = 2$ is a particularly convenient case to study from a technical perspective. It is also a natural case to study for the following descriptive reason: it has a natural extension to an arbitrary number of agents as we show below.

Given a distribution of posteriors in the population $H \in \Delta_n([0, 1])$, we can measure polarization via $\mathrm{Var}(H)$. Note that in the special case of $n = 2$ with posteriors x_1, x_2 the variance is $\frac{(x_1-x_2)^2}{4}$, which (up to a factor of 4) is exactly the case of $\alpha = 2$.

The results of [Dubins and Pitman, 1980] for $\alpha = 1$ and of [Arieli et al., 2020] for $\alpha \leq 2$ show that for the case of two agents the information structure that maximizes polarization is the one that reveals the state to one agent and provides no information to the other agent.[4] This information structure achieves an expected variance of $\frac{\mu(1-\mu)}{4}$ for information structures with prior $\mu \in [0, 1]$. We show that the same result is true for every even number of players: the information structure that maximizes polarization is the one that reveals the state to half of the agents and provides no information to the other half of the agents.

The polarization of G is the expected variance of the distribution of posteriors in the population. That is, for any $G \in \mathcal{G}$ let $V_G = \int_{\Delta([0,1])} \mathrm{Var}(H) \mathrm{d}G(H)$.

Proposition 2. Let $\mu \in [0, 1]$ be the prior. For every even n, the information structure that maximizes the polarization of n agents is the one that reveals the state to half of the agents and reveals no information to the other half. The expected polarization equals $\frac{\mu(1-\mu)}{4}$.

For every odd n, the maximal polarization is bounded between

$$\left(1 - \frac{1}{n^2}\right) \frac{\mu(1-\mu)}{4} \leq \max_{G \in \mathcal{G}_n} V_G \leq \frac{\mu(1-\mu)}{4}.$$

[4] It remains an open problem to specify the information structure that maximizes polarization for $\alpha > 2$, even for the case of two agents.

For an odd number of agents, Proposition 2 provides asymptotically tight bounds when $n \to \infty$. It remains an open problem to provide the precise characterization of polarization for an odd number of agents. Interestingly, the idea of sending full information or no information to each one of the agents is no longer optimal for an odd number of agents as demonstrated in the following example.

Example 1. For $n = 3$ and prior $\mu = \frac{1}{2}$, consider the information structure where the state is revealed to agent 3, and agents $i = 1, 2$ receive the signals s_1^i, s_2^i according to the following distribution conditional on a state:

$\omega = 0$

	s_1^2	s_2^2
s_1^1	1/3	1/3
s_2^1	1/3	0

$\omega = 1$

	s_1^2	s_2^1
s_1^1	0	1/3
s_2^1	1/3	1/3

Agent i's posterior after observing s_1^i is 1/3; her posterior after observing s_2^i is 2/3. The resulting distribution over posterior profiles is uniform over the profiles:

$$\{(\frac{1}{3}, \frac{1}{3}, 0), (\frac{1}{3}, \frac{2}{3}, 0), (\frac{2}{3}, \frac{1}{3}, 0), (\frac{1}{3}, \frac{2}{3}, 1), (\frac{2}{3}, \frac{1}{3}, 1), (\frac{2}{3}, \frac{2}{3}, 1)\}.$$

The expected variance is $\frac{2}{6}\frac{2}{81} + \frac{4}{6}\frac{2}{27} = \frac{14}{243}$, which is above $\frac{1}{18}$, where $\frac{1}{18}$ is the maximal expected variance that can be obtained by sending to every player either full information or no information about the state.

Proof of Proposition 2. First we show that $V_G \leq \frac{\mu(1-\mu)}{4}$ for every number of agents n (odd or even). This statement can be equivalently stated with respect to a feasible population's empirical posteriors. We denote $V_{\mathbf{P}} = \int_{\Delta([0,1])} \mathrm{Var}(H) \mathrm{d}\mathbf{P}(H)$ and we trivially have $\max_{G \in \mathcal{G}_n} V_G = \max_{\mathbf{P} \in \mathcal{P}_n(\mu)} V_{\mathbf{P}}$.

Let $\mathbf{P} \in \mathcal{P}_n(\mu)$. By Theorem 1 we can decompose $\mathbf{P} = \mu_{\mathbf{P}} \mathbf{P}_1 + (1 - \mu_{\mathbf{P}}) \mathbf{P}_0$ where $\mathbb{E}[\mathbf{P}^\omega] = \tilde{\mathbf{P}}^\omega$. Therefore, $V_{\mathbf{P}} = \mu_{\mathbf{P}} V_{\mathbf{P}_1} + (1 - \mu_{\mathbf{P}}) V_{\mathbf{P}_0}$. Let $\mathbf{P}' = \mu_{\mathbf{P}} \delta_{\tilde{\mathbf{P}}_1} + (1 - \mu_{\mathbf{P}}) \delta_{\tilde{\mathbf{P}}_0}$. Note that since $\mathbf{P} \in \mathcal{P}(\mu)$ it also holds that $\mathbf{P}' \in \mathcal{P}(\mu)$. We claim that $V_{\mathbf{P}} \leq V_{\mathbf{P}'}$. To see this,

note that for any $\omega = 0, 1$,

$$V_{\mathbf{P}^\omega} = \int_{\Delta([0,1])} \mathrm{Var}(H) \mathrm{d}\mathbf{P}^\omega(H) = \int_{\Delta([0,1])} \mathbb{E}_H[x^2] - (\mathbb{E}_H[x])^2 \mathrm{d}\mathbf{P}^\omega(H) = \tag{4}$$

$$\int_{\Delta([0,1])} \mathbb{E}_H[x^2] \mathrm{d}\mathbf{P}^\omega(H) - \int_{\Delta([0,1])} (\mathbb{E}_H[x])^2 \mathrm{d}\mathbf{P}^\omega(H) =$$

$$\mathbb{E}_{\tilde{\mathbf{P}}^\omega}[x^2] - \int_{\Delta([0,1])} (\mathbb{E}_H[x])^2 \mathrm{d}\mathbf{P}(H) \leq \tag{5}$$

$$\mathbb{E}_{\tilde{\mathbf{P}}^\omega}[x^2] - (\mathbb{E}_{\tilde{\mathbf{P}}^\omega}[x]))^2 = \mathrm{Var}(\delta_{\tilde{\mathbf{P}}^\omega}) \tag{6}$$

Note that equation (4) follows from $\mathrm{Var}(H) = \mathbb{E}_H[x^2] - (\mathbb{E}_H[x])^2$. Equation (5) follows from the fact that the operation $\mathbb{E}_H(x^2) : \Delta(\Delta([0,1])) \to \mathbb{R}$ is a linear operator and so for every $\mathbf{Q} \in \Delta(\Delta([0,1]))$ we have that $\int_{\Delta([0,1])} \mathbb{E}_H[x^2] \mathrm{d}\mathbf{Q}(H) = \mathbb{E}_{\tilde{\mathbf{Q}}}[x^2]$. The inequality in (6) follows from Jensen's inequality. Since $V_{\mathbf{P}'} = \mu_{\mathbf{P}} V_{\delta_{\tilde{\mathbf{P}}_1}} + (1 - \mu_{\mathbf{P}}) V_{\delta_{\mathbf{P}_0}}$ it follows that $V_{\mathbf{P}'} \geq V_{\mathbf{P}}$.

Note that $\tilde{\mathbf{P}}_1, \tilde{\mathbf{P}}_0$ satisfy

$$\mathrm{d}\tilde{\mathbf{P}}_1 = \frac{x}{\mu} \mathrm{d}\tilde{\mathbf{P}} \text{ and } \mathrm{d}\tilde{\mathbf{P}}_0 = \frac{1-x}{1-\mu} d\tilde{\mathbf{P}},$$

for some $\tilde{\mathbf{P}} \in \Delta([0,1])$ with expectation μ, and hence the variances can be written as

$$\mathrm{Var}(\tilde{\mathbf{P}}_1) = \int_{[0,1]} \frac{x^3}{\mu} \mathrm{d}\tilde{\mathbf{P}}(x) - \left(\int_{[0,1]} \frac{x^2}{\mu} \mathrm{d}\tilde{\mathbf{P}}(x) \right)^2$$

$$\mathrm{Var}(\tilde{\mathbf{P}}_0) = \int_{[0,1]} \frac{x^2(1-x)}{1-\mu} \mathrm{d}\tilde{\mathbf{P}}(x) - \left(\int_{[0,1]} \frac{x(1-x)}{1-\mu} \mathrm{d}\tilde{\mathbf{P}}(x) \right)^2$$

Therefore,

$$V_{\mathbf{P}} \leq \mu \,\mathrm{Var}(\tilde{\mathbf{P}}_1) + (1-\mu) \,\mathrm{Var}(\tilde{\mathbf{P}}_0) =$$

$$\int_{[0,1]} x^3 + x^2(1-x) \mathrm{d}\tilde{\mathbf{P}}(x) - \frac{1}{\mu} \left(\int_{[0,1]} x^2 \mathrm{d}\tilde{\mathbf{P}}(x) \right)^2 - \frac{1}{1-\mu} \left(\int_{[0,1]} x(1-x) \mathrm{d}\tilde{\mathbf{P}}(x) \right)^2 =$$

$$\int_{[0,1]} x^2 \mathrm{d}\tilde{\mathbf{P}}(x) - \frac{1}{\mu} \left(\int_{[0,1]} x^2 \mathrm{d}\tilde{\mathbf{P}}(x) \right)^2 - \frac{1}{1-\mu} \left(\mu - \int_{[0,1]} x^2 \mathrm{d}\tilde{\mathbf{P}}(x) \right)^2 =$$

$$(\frac{1+\mu}{1-\mu}) \int_{[0,1]} x^2 \mathrm{d}\tilde{\mathbf{P}}(x) - \frac{1}{\mu(1-\mu)} \left(\int_{[0,1]} x^2 \mathrm{d}\tilde{\mathbf{P}}(x) \right)^2 - \frac{\mu^2}{1-\mu}.$$

Denote $A = \int_{[0,1]} x^2 \mathrm{d}\tilde{\mathbf{P}}(x)$ and we obtain $V_{\mathbf{P}} \leq (\frac{1+\mu}{1-\mu})A - \frac{1}{\mu(1-\mu)}A^2 - \frac{\mu^2}{1-\mu}$. This is a degree 2 polynomial whose global maximum is $\frac{\mu(1-\mu)}{4}$ that is obtained at $A = \frac{\mu(1+\mu)}{2}$.

To show that $\frac{\mu(1-\mu)}{4}$ is achievable for an even number of players, we calculate the expected variance for the information structure that reveals the state to $\frac{n}{2}$ of the players and reveals nothing to the remaining $\frac{n}{2}$ players and get precisely $\frac{\mu(1-\mu)}{4}$.

To show that $(1 - \frac{1}{n^2})\frac{\mu(1-\mu)}{4}$ is achievable for an odd number of players, we calculate the expected variance for the information structure that reveals the state to $\frac{n+1}{2}$ of the players and reveals no information to the remaining $\frac{n-1}{2}$ players and get the desired expression. □

4.1 Multiple States

For an arbitrary finite state space Ω, one might measure polarization of a distribution of posteriors in the population via $\mathrm{Pol}(H) = \mathbb{E}_{x\sim H}[\|x - \mathbb{E}[H]\|_2^2]$, where x and $\mathbb{E}[H]$ are $|\Omega|$-dimensional vectors of beliefs in the simplex. Note that $\mathrm{Pol}(H) = \sum_{\omega\in\Omega} \mathrm{Var}(X_\omega)$, where X_ω denotes the random variable of the posterior probability that is assigned to state $\omega \in \Omega$. Therefore, for an arbitrary state space we obtain the following corollary.

Corollary 2. For every prior $\mu \in \Delta(\Omega)$ and every even n the information structure that maximizes the expected polarization (i.e., $\mathbb{E}_{H\sim P} \mathrm{Pol}(H)$) of n agents is the one that reveals the state to half of the agents and reveals no information to the other half.

The corollary simply follows from the fact that $\mathrm{Pol}(H) = \sum_{\omega\in\Omega} \mathrm{Var}(X_\omega)$ and the described information structure simultaneously maximizes $\mathrm{Var}(X_\omega)$ for all $\omega \in \Omega$.

5 Vectorial Feasibility

In this section we connect our results on the feasibility of distributions over *empirical distributions of beliefs* with the notion of the feasibility of distributions over *vectors of beliefs*. As mentioned in the Introduction, the latter notion of feasibility has been previously studied in [Dawid et al., 1995, Arieli et al., 2021, He et al., 2021]. To distinguish between the two notions we refer to the latter notion as *vectorial feasibility*.

The model is similar to the one presented in Section 3. Given an information structure $I \in \Delta(\Omega \times S_1 \times ... \times S_n)$, every realization of signals $(s_i)_{i\in[n]}$ induces a vector of beliefs

$(x_i(s_i))_{i\in[n]}$. Hence, every information structure induces a distribution over vectors of beliefs $\mathbf{V}_I \in \Delta((\Delta(\Omega)^n)$.

Definition 2. A distribution $\mathbf{V} \in \Delta((\Delta(\Omega)^n)$ is *vector feasible* if there exists I such that $\mathbf{V} = \mathbf{V}_I$.

Every distribution $\mathbf{V} \in \Delta((\Delta(\Omega)^n)$ induces a distribution $\mathbf{P}(\mathbf{V}) \in \Delta(\Delta_n(\Delta(\Omega)))$ simply because every vector of beliefs induces an empirical distribution of beliefs in the population. We observe that the feasibility of $\mathbf{P}(\mathbf{V})$ is a necessary but insufficient condition for the feasibility of $\mathbf{V}$. Indeed this condition is necessary, because if there exists no I that will induce the distribution of empirical distributions $\mathbf{P}(\mathbf{V})$ then $\mathbf{V}$ cannot be implemented by any I. It is also not hard to see that the feasibility of $\mathbf{P}(\mathbf{V})$ is an insufficient condition for the feasibility of[5] $\mathbf{V}$.

Therefore, the tools developed in this paper might be useful for determining the vectorial infeasibility of $\mathbf{V}$ by checking our condition for the feasibility of $\mathbf{P}(\mathbf{V})$. The following subsection demonstrates that in sufficiently symmetric cases the feasibility of $\mathbf{P}(\mathbf{V})$ serves as a necessary *and sufficient* condition for the vectorial feasibility of $\mathbf{V}$. In particular, using our results we are able to deduce new insights about feasible and infeasible vector distributions $\mathbf{V}$.

5.1 Private private information

A special case of distributions $\mathbf{V}$ whose vectorial feasibility has been studied in [He et al., 2021, Arieli et al., 2021, Gutmann et al., 1991b] is that of $|\Omega| = 2$ and a *product* distribution $V = \Pi_{i\in[n]} Q_i$ where $Q_i \in \Delta([0,1])$ is the distribution of the beliefs of agent i. [He et al., 2021] refer to this class of distributions as *private private information* because the private information of agents other than i does not affect their beliefs about agent i's belief. In any case, the belief of every agent $j \neq i$ about agent i's belief will be Q_i.

[5] To construct a counterexample consider $|\Omega| = 2$, $n = 2$ and consider any *feasible* vector distribution $\mathbf{V}'$ with $\mathbf{V}'(a,b) > 0$ for some $a \neq b$. Let $\mathbf{V}$ be a distribution that is obtained from $\mathbf{V}'$ by moving some positive mass from (a,b) to (b,a). The distribution $\mathbf{V}$ becomes vector infeasible because it violates the martingale condition for both agents. However $\mathbf{P}(\mathbf{V}) = \mathbf{P}(\mathbf{V}')$ remains feasible.

Another counterexample that does not violate the martingale condition for any player is the following. Let $|\Omega| = 2$ and $n = 3$. The distribution $\mathbf{V} = \frac{1}{2}\delta_{(\frac{1}{3},\frac{1}{3},\frac{2}{3})} + \frac{1}{2}\delta_{(\frac{2}{3},\frac{2}{3},\frac{1}{3})}$ is vector infeasible because agents 1 and 2 are agreeing to disagree with agent 3 in both points of the support of this distribution. However, $\mathbf{P}(\mathbf{V})$, which is the distribution where with probability $\frac{1}{2}$ two of the three agents have a posterior of $\frac{j}{3}$ (without specifying their identity) and the remaining agent has a posterior of $\frac{3-j}{3}$ for $j = 1, 2$, is feasible by Corollary 1.

For a discussion on the attractiveness of this property see [He et al., 2021, Arieli et al., 2021]. A special case of private private information that has received attention in the existing literature (see [He et al., 2021, Arieli et al., 2021, Gutmann et al., 1991b]) is the *symmetric* case where $V = Q^n$. The characterization of the feasible and infeasible $V = Q^n$ is known for the case of $n = 2$ ([He et al., 2021]) but remains unclear for $n \geq 3$.

Consider now the case of a finitely supported Q. For the case of $V = Q^n$ the induced distribution over empirical distributions $\mathbf{P}(\mathbf{V})$ is the multinomial distribution. More formally, assume that Q is supported on $x_1, ..., x_k \in [0, 1]$ and assigns a probability of α_i to x_i. Let $\mathrm{MN}(n, Q)$ be the multinomial distribution that draws n samples from $\{x_i\}_{i\in[k]}$ according to Q. In the standard definition of the multinomial distribution it is assumed that the outcome is a vector of integers $(n_1, ..., n_k)$ that sum up to n and indicates the number of times $x_1, ..., x_k$ have been chosen. In our case, we refer to the outcome of $\mathrm{MN}(n, Q)$ as $(\frac{n_1}{n}, ..., \frac{n_k}{n})$ and hence the outcome is an element of $\Delta_n(\Delta(\Omega))$. The following proposition shows that in this case the feasibility of $\mathbf{P}(\mathbf{V})$ is a necessary and sufficient condition for the vectorial feasibility of $\mathbf{V} = Q^n$. In particular, it provides a novel characterization of vectorial feasibility for this class of distributions.

Proposition 3. For a binary state space Ω and $Q \in \Delta([0, 1])$ the product distribution $\mathbf{V} = Q^n$ for n agents is vector feasible if and only if the multinomial distribution $\mathrm{MN}(n, Q)$ is a mean-preserving spread of[6] $(1 - \mu)Q^{\omega=0} + \mu Q^{\omega=1}$.

Proof. Denote $\mathbf{P} = \mathbf{P}(\mathbf{V}) = \mathrm{MN}(n, Q)$. Note that $\tilde{\mathbf{P}} = Q$. By Theorem 1, $\mathbf{P}$ is feasible if and only if $\mathbf{P}$ is a mean-preserving spread of $(1 - \mu)Q^{\omega=0} + \mu Q^{\omega=1}$.

If $\mathbf{P} = \mathbf{P}(\mathbf{V})$ is infeasible then $\mathbf{V}$ is not vector feasible, as discussed above.

If $\mathbf{P}$ is feasible, we observe that $\mathbf{V}$ can be viewed as the *unique* distribution $\mathbf{W}$ with the following two properties:

1. $\mathbf{P}(\mathbf{W}) = \mathbf{P}$.
2. $\mathbf{W}$ assigns an equal probability to every two vectors $(y_1, ..., y_n), (y'_1, ..., y'_n) \in \{x_1, ..., x_k\}^n$ with the same empirical distributions of beliefs.

It is easy to check that our construction of the information structure in the proof of

[6] We recall that $Q^\omega \in \Delta([0, 1])$ denotes the distribution of beliefs conditional on state ω in the case where the ex-ante distribution of beliefs is Q.

Theorem 1 induces a distribution over vectors of beliefs that satisfies these two properties and hence this information structure implements $\mathbf{V}$. Thus, $\mathbf{V}$ is feasible. $\square$

Determining whether $\mathrm{MN}(n, Q)$ is a mean-preserving spread of $(1-\mu)Q_0+\mu Q_1$ might not be an easy task in the case where k (i.e., the support size of Q) is large. This is because $\mathrm{MN}(n, Q)$ is a distribution over elements in $\Delta([k])$, and in high dimensions the mean-preserving condition might not be easy to check. However, for $k = 2$ the multinomial distribution becomes simply a binomial distribution and the mean-preserving condition is easily verifiable by Lemma 1. Thus, for Q with a binary support we get the following condition for the vectorial feasibility of $\mathbf{V} = Q^n$.

Corollary 3. Let $Q \in \Delta(\{a, b\})$ where $0 < a < \mu < b < 1$. The distribution $\mathbf{V} = Q^n$ is vector feasible (for n agents) if and only if[7] $\mathbb{E}[\mathrm{Bin}(n, \frac{\mu-a}{b-a})|_\mu] \leq \frac{(\mu-a)b}{(b-a)\mu}$.

That is, the corollary states that in order to determine the feasibility of $\mathbf{V} = Q^n$ one should calculate the expectation of the μ-quantile of the Binomial distribution. If it exceeds some constant, then the distribution is infeasible. Otherwise it is feasible. Below is a case where we can provide an explicit formula for the feasibility for every number of agents.

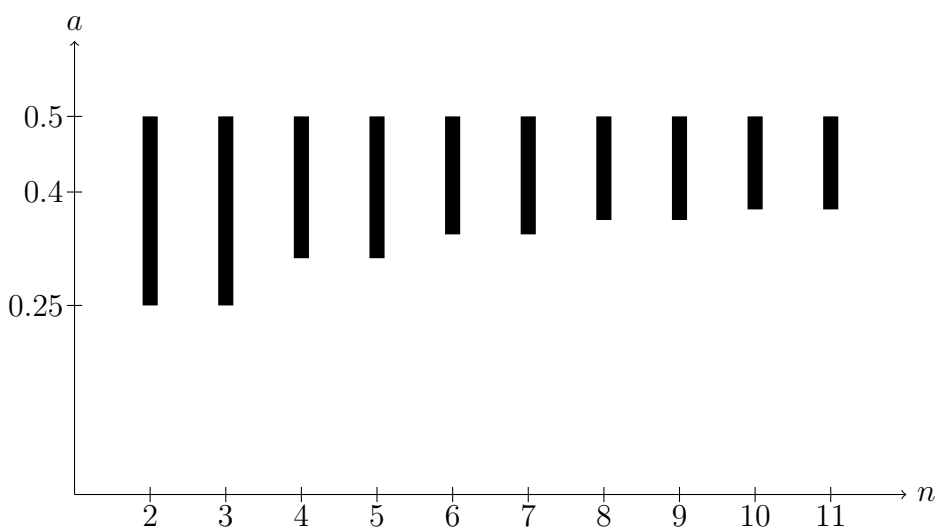

Figure 1: The values of a for which the product distribution Q^n is feasible as a function of the number of agents n.

[7] We adopt here our previous convention for the multinomial distribution, and assume that the Binomial distribution receives values in $\{0, \frac{1}{n}, ..., 1\}$ and not in $\{0, 1, ..., n\}$.

Example 2. Let $\mu = \frac{1}{2}$ and let $Q = \frac{1}{2}\delta_a + \frac{1}{2}\delta_{1-a}$ for $0 < a < \frac{1}{2}$. For $n = 2m$ and for $n = 2m+1$, the distribution $\mathbf{V} = Q^n$ is feasible if and only if $a \geq \frac{1}{2} - \binom{2m}{m}2^{-2m-1}$ (see Figure 1). This observation is a consequence of Corollary 3, where the expectation of the $\frac{1}{2}$-quantile of the $\mathrm{Bin}(n, \frac{1}{2})$ distribution can be explicitly computed and is provided in the following lemma.

Lemma 2. For every $m \in \mathbb{N}$ we have

$$\mathbb{E}[\mathrm{Bin}(2m, p)|_{\frac{1}{2}}] = \mathbb{E}[\mathrm{Bin}(2m+1, p)|_{\frac{1}{2}}] = \frac{1}{2} - \binom{2m}{m}2^{-2m-1}.$$

The proof of this lemma is relegated to Appendix B.

Asymptotically, for large population of n agents, the threshold value of a for which $V = Q^n$ is feasible is $a \geq \frac{1}{2} - \binom{2m}{m}2^{-2m-1} \approx \frac{1}{2} - \frac{1}{\sqrt{2\pi n}}$. Another interesting phenomenon is that the threshold value of a is identical for an even n and for $n+1$.

6 Private Bayesian Persuasion of a Homogeneous Population

We next consider an application of our result to private Bayesian persuasion. As in [Arieli and Babichenko, 2019], consider a sender who is a marketer that tries to sell a product to a group of agents. Consider a binary state space $\Omega = \{0, 1\}$ with a common prior μ. The agents' action set is $A = \{0, 1\}$, where action $a = 1$ represents a decision to buy the product and action $a = 0$ represents a decision not to buy the product. Playing $a = 0$ yields a utility of 0. The state represents the quality of the product. In $\omega = 1$ the product is of good quality and worth 1 to the agents while in state $\omega = 0$ it is worth $-\frac{\tau}{1-\tau}$ to the agent, where $\tau \in (0, 1)$. Thus, at a posterior of τ the agents are indifferent between buying and not buying the product. We assume that the agents share no payoff externalities. Let $u : [0, 1] \to \mathbb{R}$ be the utility of the sender, where $u(x)$ represents her utility from persuading a fraction x of agents to take action 1. We assume that u is non-decreasing and that $\mu < \tau$ (otherwise the policy that reveals no information would be optimal). We denote by V_n the optimal utility the sender can achieve in an interaction with n agents.

[Arieli and Babichenko, 2019, Babichenko and Barman, 2016] study a heterogeneous variant of this setting, where each agent possesses an individual threshold $\tau_i \in (0,1)$. In the heterogeneous case a clean formula for V_n has been deduced only in the case where u is concave; see [Arieli and Babichenko, 2019].[8] Here we observe that the homogeneous population variant of the problem can be solved using our main result (Theorem 1) and a clean formula for V_n can be provided.

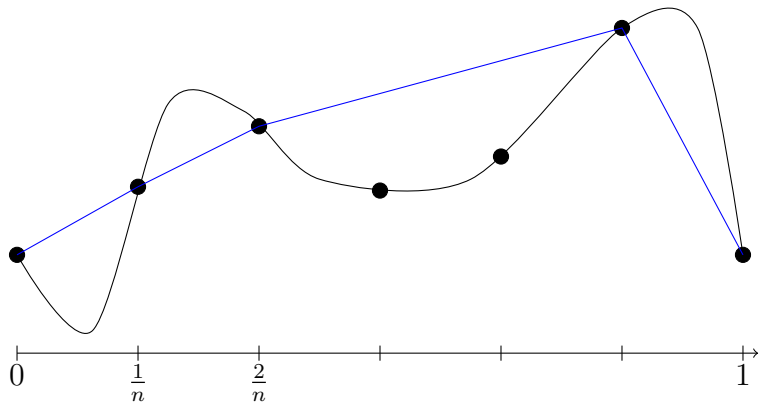

Figure 2: Concavification over a grid. The function u is displayed in black. The concavification over the grid $\mathrm{cav}_n(u)$ is displayed in blue.

We denote by $\mathrm{cav}_n(u) : [0,1] \to \mathbb{R}$ the concavification of u when u is restricted to the $\frac{1}{n}$-grid points $\{0, \frac{1}{n}, \frac{2}{n}, ..., 1\}$ only (see Figure 2). Namely, $\mathrm{cav}_n(u)$ is the minimial concave function that satisfies $\mathrm{cav}_n(u)(x) \geq u(x)$ for every $x \in \{0, \frac{1}{n}, \frac{2}{n}, ..., 1\}$. Equivalently, for every $y \in [0,1]$ we define

$$\mathrm{cav}_n(u)(y) = \max_{Q \in \Delta_n(\{0,1\}), \mathbb{E}[Q]=y} \mathbb{E}_{x \sim Q}[u(x)]. \tag{7}$$

The characterization of the sender's optimal value is as follows.

Proposition 4. The optimal utility for the sender is

$$V_n = \mu u(1) + (1-\mu)\,\mathrm{cav}_n(u)\left(\frac{\mu(1-\tau)}{\tau(1-\mu)}\right).$$

The formula for the value can be interpreted as follows. In state $\omega = 1$ all agents take action 1 and the utility is $u(1)$. In state $\omega = 0$ an expected fraction of $\frac{\mu(1-\tau)}{\tau(1-\mu)}$ agents adopt the product and the sender is allowed to optimize over all the distributions with

[8] A polynomial algorithm for computing V_n has been deduced in [Babichenko and Barman, 2016].

expectation $\frac{\mu(1-\tau)}{\tau(1-\mu)}$ in order to achieve $\text{cav}_n(u)(\frac{\mu(1-\tau)}{\tau(1-\mu)})$.

Note that the concavification in Proposition 4 differs from the classic characterization of the value via concavification ([Kamenica and Gentzkow, 2011]). The standard characterization applies concavification to a function whose domain is that of the posterior beliefs. Our characterization (for this special case) applies concavification to a function whose domain is that of a fraction of the buyers.

Proposition 4 also provides an asymptotic characterization of the value for large populations. Denote $V^* = \lim_{n\to\infty} V_n$ and denote by $\text{cav}(u)$ the standard notion of concavification (without restricting it to a grid). By Proposition 4 we have

$$V^* = \mu u(1) + (1-\mu)\,\text{cav}(u)(\frac{\mu(1-\tau)}{\tau(1-\mu)})$$

because $\text{cav}_n(u)$ point-wise converges to $\text{cav}(u)$ as $n \to \infty$ for non-decreasing functions u.

Proof of Proposition 4. The idea is to utilize Corollary 1. To do so, we should be able to identify the possible posteriors that agents could have in an optimal policy, and to prove that there are only two possible values for these posteriors. To this end, we utilize the following observation from [Arieli and Babichenko, 2019].

Lemma 3 ([Arieli and Babichenko, 2019])**.** There exists an optimal policy with the following properties:

- The policy uses binary signals $S_i = \{s_0^i, s_1^i\}$ for each agent i, and each agent takes action 1 if and only if she receives the signal s_1^i.
- In state $\omega = 1$ signal s_1^i is sent to all agents with probability 1.
- In state $\omega = 0$ each agent i gets the signal s_1^i with probability[9] $\frac{\mu(1-\tau)}{\tau(1-\mu)}$.

By Lemma 3 we deduce that the only two possible posteriors for an agent are 0 and τ. Now we can use Corollary 1, which states that $\mathbf{P} \in \Delta(\{0, \frac{1}{n}, \frac{2}{n}, ..., 1\})$ is feasible if and only if $\mathbf{P}$ is a mean-preserving spread of $\mu\delta_1 + (1-\mu)\delta_{\frac{\mu(1-\tau)}{\tau(1-\mu)}}$, where $\mathbf{P}$ is the number of agents with posterior τ. This set of mean-preserving spreads consists of all

[9] The Lemma does not indicate on the correlation of the s_1^i signals among the agents conditional on state $\omega = 0$.

$\mu\delta_1 + (1-\mu)Q$ where $Q \in \Delta(\{0, \frac{1}{n}, \frac{2}{n}, ..., 1\})$ satisfy $\mathbb{E}[Q] = \frac{\mu(1-\tau)}{\tau(1-\mu)}$ and the utility of the sender for each such mean-preserving spread is $\mu u(1) + (1-\mu)\mathbb{E}_{x\sim Q}[u(x)]$, which is precisely the concavification formula of Equation (7). □

References

Itai Arieli and Yakov Babichenko. Private Bayesian persuasion. *Journal of Economic Theory*, 182:185–217, 2019.

Itai Arieli, Yakov Babichenko, and Rann Smorodinsky. Identifiable information structures. *Games and Economic Behavior*, 120:16–27, 2020.

Itai Arieli, Yakov Babichenko, Fedor Sandomirskiy, and Omer Tamuz. Feasible joint posterior beliefs. *Journal of Political Economy*, 129(9):2546–2594, 2021.

Robert J Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239, 1976.

Robert J Aumann, Michael Maschler, and Richard E Stearns. *Repeated Games with Incomplete Information*. MIT Press, 1995.

Yakov Babichenko and Siddharth Barman. Computational aspects of private Bayesian persuasion. *arXiv preprint arXiv:1603.01444*, 2016.

David Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102. University of California Press, 1951.

Krzysztof Burdzy and Soumik Pal. Contradictory predictions. *arXiv preprint arXiv:1912.00126*, 2019.

Krzysztof Burdzy and Jim Pitman. Bounds on the probability of radically different opinions. *Electronic Communications in Probability*, 25:1–12, 2020.

AP Dawid, MH DeGroot, J Mortera, R Cooke, S French, C Genest, MJ Schervish, DV Lindley, KJ McConway, and RL Winkler. Coherent combination of experts' opinions. *Test*, 4(2):263–313, 1995.

Henrique De Oliveira and Rohit Lamba. Rationalizing dynamic choices. *Available at SSRN 3332092*, 2019.

Lester E Dubins and Jim Pitman. A maximal inequality for skew fields. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 52(3):219–227, 1980.

Shaddin Dughmi and Haifeng Xu. Algorithmic Bayesian persuasion. *SIAM Journal on Computing*, (3):STOC16–68, 2019.

Sam Gutmann, JHB Kemperman, JA Reeds, and Larry A Shepp. Existence of probability measures with given marginals. *The Annals of Probability*, 19(4):1781–1797, 1991a.

Sam Gutmann, JHB Kemperman, JA Reeds, and Larry A Shepp. Existence of probability measures with given marginals. *The Annals of Probability*, 19(4):1781–1797, 1991b.

Kevin He, Fedor Sandomirskiy, and Omer Tamuz. Private private information. *arXiv preprint arXiv:2112.14356*, 2021.

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

Paul Milgrom and Nancy Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.

Pooya Molavi. Tests of Bayesian rationality. *arXiv preprint arXiv:2109.07007*, 2021.

Stephen E Morris. No trade and feasible joint posterior beliefs. *Working paper*, 2020.

Eran Shmaya and Leeat Yariv. Experiments on decisions under uncertainty: A theoretical framework. *American Economic Review*, 106(7):1775–1801, 2016.

A Extending the Proof of Theorem 1 to infinitely supported distributions

We have proved in Section 3.3 that a finitely supported feasible $\mathbf{P}$ must be a mean-preserving spread of $\mathbf{P}'$. Here we extend this result to arbitrary feasible distributions

$\mathbf{P}$.

Let G be an information structure G for n agents where $\{S_i\}_{i=1}^n$ are (possibly infinite) measurable sets. Note that the measure $\mathbf{P}_G$ that is induced by G can be weakly approximated using $\mathbf{P}_{G'}$ when G' is a finite information structure. To do this, we use the following standard trick. First we partition $\Delta(\Omega)$ using a finite partition $\mathcal{D}_\delta$ such that each element $D \in \mathcal{D}_\delta$ is a convex set and has a diameter at most δ. We define the new signal space $S'_i = \mathcal{D}_\delta$ for every agent i. We then couple together G and G' as follows. We draw a signal profile according to G and tell each agent i only of the partition element in $\mathcal{D}_\delta$ in which the posterior $\mathbf{P}(\omega = 1|s_i)$ lies. It is easy to see that $\mathbf{P}_{G'}$ weakly converges to $\mathbf{P}$ as δ goes to 0.

It now remains to show that if $\{\mathbf{P}_k\}_k \subset \Delta(\Delta_n(\Delta(\Omega)))$ is a sequence of measures with finite support that weakly converges to $\mathbf{P}$ such that $\mathbf{P}_k$ is a mean-preserving spread of $\mathbf{P}'$ for every k, then $\mathbf{P}$ also satisfies the this condition.

Note that $\mathbf{P}_k = \sum_{\omega\in\Omega} \mu_\omega \mathbf{P}_k^\omega$. Since $\Delta(\Delta_n(\Delta(\Omega)))$ is a weakly compact space we can assume (by taking subsequences if necessary) that all sequences $\{\mathbf{P}_k^\omega\}_k$ weakly converge to a measure $\mathbf{P}^\omega$ for every state $\omega \in \Omega$. Therefore, $\mathbf{P} = \sum_{\omega\in\Omega} \mu_\omega \mathbf{P}_\omega$. We further note that taking the expected measure is a weakly continuous operator from $\Delta(\Delta_n(\Delta(\Omega)))$ to $\Delta(\Delta(\Omega))$. Therefore, the sequences $\{\tilde{\mathbf{P}}_k\}_k$ and $\{\tilde{\mathbf{P}}_k^\omega\}_k$ converge to $\tilde{\mathbf{P}}$ and $\tilde{\mathbf{P}}^\omega$ respectively for every ω. Therefore $\mathbf{P}$ is a mean-preserving spread of $\mathbf{P}' = \sum_{\omega\in\Omega} \mu_\omega \delta_{\tilde{\mathbf{P}}^\omega}$ as desired.

Furthermore, since for every k and $\omega \in \Omega$,

$$\mathrm{d}\tilde{\mathbf{P}}_k^\omega(x) = \frac{1}{\mu_\omega} x_\omega \mathrm{d}\tilde{\mathbf{P}}_k(x),$$

it follows that in the limit for every $\omega \in \Omega$,

$$\mathrm{d}\tilde{\mathbf{P}}^\omega(x) = \frac{1}{\mu_\omega} x_\omega \mathrm{d}\tilde{\mathbf{P}}(x),$$

as desired.

B Proof of Lemma 2

The distribution $\mathrm{Bin}(2m+1, \frac{1}{2})|_{\frac{1}{2}}$ assigns a probability of $\binom{2m+1}{k}2^{-2m}$ to $\frac{k}{2m+1}$ for every $k \leq m$, hence its expectation is given by

$$\begin{aligned}
\mathbb{E}[\mathrm{Bin}(2m+1, \frac{1}{2})|_{\frac{1}{2}}] &= \sum_{k=1}^{m} \binom{2m+1}{k} 2^{-2m} \frac{k}{2m+1} = \frac{2^{-2m}}{2m+1} \sum_{k=1}^{m} \binom{2m+1}{k} k \\
&= \frac{2^{-2m}}{2m+1}(2m+1) \sum_{k=0}^{m-1} \binom{2m}{k} = 2^{-2m}\frac{1}{2}\left(2^{2m} - \binom{2m}{m}\right) \\
&= \frac{1}{2} - \binom{2m}{m} 2^{-2m-1}
\end{aligned} \tag{8}$$

where the third equality follows from the fact that both expressions $\sum_{k=1}^{m} \binom{2m+1}{k} k$ and $(2m+1)\sum_{k=0}^{m-1} \binom{2m}{k}$ count the number of possible choices of a subset of at most k people and their leader out of a group of $2m+1$ people. The first expression first chooses the group and thereafter the leader. The second expression first chooses the leader and thereafter complements it with a subset.

The distribution $\mathrm{Bin}(2m, \frac{1}{2})|_{\frac{1}{2}}$ assigns a probability of $\binom{2m}{k}2^{-2m+1}$ to $\frac{k}{2m}$ for every $k < m$ and assigns a probability of $\binom{2m}{m}2^{-2m}$ to $\frac{1}{2}$, and hence its expectation is given by

$$\begin{aligned}
\mathbb{E}[\mathrm{Bin}(2m, \frac{1}{2})|_{\frac{1}{2}}] &= \sum_{k=1}^{m-1} \binom{2m}{k} 2^{-2m+1} \frac{k}{2m} + \binom{2m}{m} 2^{-2m-1} \\
&= \frac{2^{-2m+1}}{2m} \sum_{k=1}^{m-1} \binom{2m}{k} k + \binom{2m}{m} 2^{-2m-1} \\
&= \frac{2^{-2m+1}}{2m} 2m \sum_{k=0}^{m-2} \binom{2m-1}{k} + \binom{2m}{m} 2^{-2m-1} \\
&= 2^{-2m+1}\frac{1}{2}\left(2^{2m-1} - 2\binom{2m-1}{m-1}\right) + \binom{2m}{m} 2^{-2m-1} \\
&= \frac{1}{2} - 2^{-2m+1}\binom{2m-1}{m-1} + \binom{2m}{m} 2^{-2m-1} \\
&= \frac{1}{2} - \binom{2m}{m} 2^{-2m-1}
\end{aligned}$$

where the argument for the third equality is similar to the third equality in Equation (8).