

THE INSTITUTE
OF STATISTICS

THE CONSOLIDATED UNIVERSITY
OF NORTH CAROLINA



ON THE VALUE OF INFORMATION IN
MULTI-AGENT DECISION THEORY

by

B. Bassan and M. Scarsini

March 1994

Mimeo Series #2316

DEPARTMENT OF STATISTICS
Chapel Hill, North Carolina

On the value of information in multi-agent decision theory

by

Bruno Bassan *

and

Marco Scarsini **

ABSTRACT

Several agents, each with his own opinion (probability measure), make decisions in order to maximize their expected utility. A *super partes* person, "the chief," releases information with the goal of maximizing a social expected utility, which is an increasing function of the agents' utilities. Additional bits of information are individually beneficial to each agent, but might be socially detrimental if the social utility is concave and therefore diversification is valuable.

In this paper, information is modelled by filtrations on a suitable probability space, and the problem of establishing how much information ought to be released is tackled. Two situations are examined, in which the chief either updates or does not update her opinion.

* Dipartimento di Matematica, Politecnico di Milano, P.le Leonardo da Vinci 32, I-20133 Milano, Italy, brubas@ipmma1.polimi.it

** Dipartimento di Metodi Quantitativi, Università D'Annunzio, Viale Pindaro 42, I-65127 Pescara, Italy, scarsini@itcaspur.caspur.it, on leave at Department of Statistics, University of North Carolina, Chapel Hill, NC 27599-3260, USA, scarsini@stat.unc.edu

1. Introduction.

In a village, k shepherds pasture their flocks. At some point, the grazing grounds become parched, and the chief of the village decides that all the shepherds must move to a different area. It is known that wolves dwell in all but one of the n available paths, but it is not known which path is free. Each shepherd has his own opinion (i.e. his own subjective probability) about the free path, and so does the chief. A scout can be sent by the chief to examine the paths, but the information that he reports is unreliable, namely the probability of it being correct is less than one. Nevertheless, the information provided by the scout accrues, and is beneficial for the choice of each individual shepherd, even if it never leads to certainty.

The chief knows everything relevant about the shepherds, namely their subjective probabilities, their utility functions and therefore the way information affects their decisions. The chief embodies a social utility, which is a monotone function of the sum of the shepherds' utilities. This means that *ceteris paribus* any increase in the utility of any shepherd increases the social utility, and that all shepherds are alike in the eye of the chief. Since it is socially important that at least one flock survive, the social utility function is concave. The behavior of the shepherds and of the chief is inspired by maximization of expected utility.

Information is released by the chief *coram populo*, namely every shepherd receives the same amount of information. The chief faces two contrasting tendencies as she releases more information: the more each shepherd is informed, the more likely he is to make the good decision, and this is socially desirable; on the other hand, the more the shepherds know, the less likely they are to diversify their behavior. This, in view of the concavity of the social utility function, is in general socially detrimental.

All information which is acquired is released, but the chief has the power to limit the information acquisition process at any time t by not sending out the scout anymore and forcing the shepherds to act, i.e. to choose a path, on the basis on the information reported by the scout up to time t .

The purpose of this paper is to determine the (socially) optimal amount of information to be released.

The “fairy tale” of the shepherds and the chief is taken from Erev, Wallsten and Neal (1991), whose aim is to study qualitative models of information release from a psychological viewpoint. Our mathematical formulation is suitable for some economic applications. For instance, consider the following research and development problem: the Government gives grants to several research laboratories to develop a therapy for some disease. The different laboratories act in complete autonomy, and can choose among several techniques. The Government might obtain some side information about the techniques (their expected effectiveness, expected cost, etc.). Should the Government get this information and release it? This would imply that some low-probability-of-success-techniques would be discarded, in favor of better ones; this would benefit, most likely, the individual laboratories that switch, but not the whole society, since one of the discarded techniques could actually prove successful. As a further example, think of a country which allows several oil companies to drill on its territory. The more information it releases, the more the companies will concentrate in the “most favorable” area, and this might prove very bad if the area turns out to be dry.

Kadane (1993) contains a nice survey of models with several agents having different subjective opinions. The topic we consider in this article is related to the problem of opinion merging (see e.g. Blackwell and Dubins (1962), Kalai and Lehrer (1991), Lehrer and Smorodinsky (1993)). In our framework the conditions for merging of the opinions of the shepherds to the opinion of the chief are satisfied (the measure of the chief is absolutely continuous with respect to the other measures). Our concern, though, is not the merging itself, but to stop the flow of information before the opinions merge “too much”, if this leads to a behavior that is socially dangerous, and if diversification has an intrinsic value.

We examine two different learning mechanisms for the chief: in the first one the accruing of information does not alter her opinion, in the second one it does.

Section 2 illustrates the general model and the particular example that will be dealt with in the rest of the paper: two shepherds have to select one path out of two, namely they have to predict the outcome of a binary random variable X_0 ; the flow of information will be represented by a sequence of random variables, exchangeable

among themselves and with X_0 . In Section 3 we consider the case in which the chief does not learn, and we give some sufficient conditions under which it is optimal to release all the available information. Typically, these conditions will depend on the degree of concavity of the social utility function. In Section 4 the chief will be allowed to update her opinion.

2. The model.

Let us fix first some notation. The measurable space in the background will be denoted by (Ω, \mathcal{F}) , and $\mathcal{P}(\Omega, \mathcal{F})$ will denote the space of all σ -additive probability measures on it. We shall consider $k + 1$ probability measures $P_0, P_1, \dots, P_k \in \mathcal{P}(\Omega, \mathcal{F})$, representing the opinions of the chief of the village (P_0) and of k shepherds (P_1, \dots, P_k). The space (Ω, \mathcal{F}) will be endowed with a filtration $\{\mathcal{F}_t \mid t \in T\}$, and $P_j^t(\cdot)$ ($j = 0, 1, \dots, k$) will denote a regular version of the conditional probability $P_j(\cdot \mid \mathcal{F}_t)$. The parameter t will be interpreted as time and, to keep things simple, will be assumed to be discrete (e.g. the number of times the scout goes out to get information), $T = \mathbb{N} \cup \{\infty\}$.

In order to describe the behavior of the shepherds, we introduce a space \mathcal{A} of actions, k utility functions $u_1, \dots, u_k : \mathcal{A} \times \Omega \rightarrow \mathbb{R}$ and k decision functions $d_1, \dots, d_k : \mathcal{P}(\Omega, \mathcal{F}) \rightarrow \mathcal{A}$, to be defined below.

The space of bounded utility functions $u : \mathcal{A} \times \Omega \rightarrow \mathbb{R}$ will be denoted by \mathcal{U} . Without loss of generality, we may assume that $0 \leq u \leq 1 \forall u \in \mathcal{U}$.

Given a measure $P \in \mathcal{P}(\Omega, \mathcal{F})$ and an integrable function f , we write

$$P[f] = \int_{\Omega} f dP.$$

The decision of the j -th shepherd will be to adopt the action that maximizes his expected utility $P_j[u_j]$, namely

$$d_j(P_j) = \arg \max_{a \in \mathcal{A}} P_j[u_j(a, \cdot)].$$

Conditioning on the information \mathcal{F}_t , we can define the (random) optimal decision d_j^t by

$$d_j^t(P_j, \omega) = d_j(P_j^t) = \arg \max_{a \in \mathcal{A}} P_j[u_j(a, \cdot) \mid \mathcal{F}_t]. \quad (2.1)$$

The maximum expected utility of the j -th shepherd at time t is given by

$$P_j [u_j(d_j^t(P_j), \cdot) \mid \mathcal{F}_t] = \max_{a \in \mathcal{A}} P_j [u_j(a, \cdot) \mid \mathcal{F}_t] = \max_{a \in \mathcal{A}} P_j^t [u_j(a, \cdot)].$$

We assume that the conditions for the existence of the above maxima are satisfied.

The social utility will be personified by the chief of the village, with utility $u_0 : \mathcal{U}^k \times \mathcal{A}^k \times \Omega \rightarrow \mathbb{R}$ given by

$$u_0(u_1, \dots, u_k, a_1, \dots, a_k, \omega) = g\left(\sum_{j=1}^k u_j(a_j, \omega)\right),$$

where $g : \mathbb{R} \rightarrow \mathbb{R}$ is increasing and bounded.

The maximum conditional expected social utility is represented by a function $W : \mathcal{U}^k \times \mathcal{P}^k \times T \times \Omega \rightarrow \mathbb{R}$ given by

$$\begin{aligned} W(u_1, \dots, u_k, P_1, \dots, P_k, t, \omega) &= P_0 [u_0(u_1, \dots, u_k, d_1^t(P_1), \dots, d_k^t(P_k), \cdot) \mid \mathcal{F}_t] \\ &= P_0 \left[g\left(\sum_{j=1}^k u_j(d_j^t(P_j), \cdot)\right) \mid \mathcal{F}_t \right]. \end{aligned}$$

The aim of the chief is to maximize, with respect to t her (social) expected utility.

The chief will make her decision about how much information to release (namely, will pick an “optimal” t) once and for all at time 0. Thus her goal is to find

$$\arg \max_{t \in T} P_0 [W(u_1, \dots, u_k, P_1, \dots, P_k, t, \cdot)] = \arg \max_{t \in T} P_0 \left[g\left(\sum_{j=1}^k u_j(d_j^t(P_j), \cdot)\right) \right].$$

In the sequel, for the sake of simplicity, we shall consider a model with only $k = 2$ shepherds, who must choose which of the two available narrow paths, labeled 0 and 1, they should go through with their sheep, knowing that in one of them there are wolves. The action space is formed by the two points b_0 and b_1 , denoting which path they choose to follow. On Ω , we define a random variable X_0 which takes the value 1 if path 1 is free of wolves, and 0 if path 0 is free of wolves. At each time t the scout tells the chief where he thinks that the wolves are, i.e. he makes known to her the value of the random variable X_t , which has the same distribution as X_0 .

The fact that the scout yields valuable information is modelled by assuming that X_0, X_1, X_2, \dots is an exchangeable family. The filtration $\{\mathcal{F}_t \mid t \in T\}$ is the natural filtration for the process $\{X_t \mid t = 1, 2, \dots\}$:

$$\mathcal{F}_t = \sigma(\{X_1, X_2, \dots, X_t\}), \quad \mathcal{F}_\infty = \cup_t \mathcal{F}_t.$$

Notice that X_0 is not assumed to be \mathcal{F}_t measurable, $\forall t \in T$.

In fact, if X_0 were \mathcal{F}_t -measurable, for some $t \in T$, then the problem would be trivial: Releasing as much information as possible would lead to certainty, and therefore would be optimal. The exchangeability assumption allows us to avoid this trivial case, and at the same time allows for dependence among the random variables involved, which is necessary for the information to be influential.

In order to shed light on the above assumptions, consider the following example. The life or death of an agent depends on the outcome of a toss of a certain coin. He is entitled to perform as many trial tosses of that coin as he wants in order to choose heads or tails, before the “true” toss. The tosses are not i.i.d., but exchangeable.

The utilities of the two shepherds represent a death-life situation, namely

$$u_1(b_i, \omega) = u_2(b_i, \omega) = \begin{cases} 1 & \text{if } X_0(\omega) = i \\ 0 & \text{otherwise} \end{cases} \quad i = 0, 1. \quad (2.2)$$

The utility of the chief, in turn, is

$$u_0(u_1, u_2, a_1, a_2, \omega) = \begin{cases} 1 & \text{if } u_1(a_1, \omega) + u_2(a_2, \omega) = 2 \\ \theta & \text{if } u_1(a_1, \omega) + u_2(a_2, \omega) = 1 \\ 0 & \text{if } u_1(a_1, \omega) + u_2(a_2, \omega) = 0 \end{cases} \quad (2.3)$$

where $\theta \in [0, 1]$. The parameter θ can be interpreted as an index of her risk-aversion in the Arrow-Pratt sense: the higher θ , the more she values the fact that at least one flock survives.

The fact that the shepherds consider the infinite sequence $\{X_n \mid n \in \mathbb{N}\}$ exchangeable implies, by de Finetti’s Theorem (de Finetti (1930,1937)) that $\sum_{i=1}^t X_i / t$ converges almost surely and in r -th mean ($r \in \mathbb{N}$) to a random variable Λ , and, conditionally on Λ , $\{X_n \mid n \in \mathbb{N}\}$ is an i.i.d. Bernoulli sequence, with probability of success equal to Λ , namely

$$P_j[X_0 \mid \sigma(\Lambda)] = \Lambda \text{ a.s. } j = 1, 2. \quad (2.4)$$

If furthermore we assume that, for all t , $P_j[X_0|\mathcal{F}_t]$ is an arithmetic mean of the prior opinion $P_j[X_0]$ and the observed frequency $\sum_{i=1}^t X_i/t$, i.e.

$$P_j[X_0|\mathcal{F}_t] = \gamma_t^{(j)} P_j[X_0] + (1 - \gamma_t^{(j)}) \frac{\sum_{i=1}^t X_i}{t}, \quad (2.5)$$

then, by a result of Diaconis and Ylvisaker (1979), according to the j -th shepherd the random variable Λ has a beta distribution with parameters (α_j, β_j) , where

$$P_j[X_0] = \frac{\alpha_j}{\alpha_j + \beta_j}$$

$$\gamma_t^{(j)} = \frac{\alpha_j + \beta_j}{\alpha_j + \beta_j + t}.$$

This particular form of updating opinions has an intuitive appeal and allows simplified computations, but it is not really crucial in the analysis. Less regular prior distributions for Λ lead to qualitatively similar results.

In this particular setting, the model described in Section 2 reads as follows: First of all, since the beta is conjugate for the Bernoulli model, conditionally on \mathcal{F}_t the P_j -distribution of Λ is again a beta with parameters $(\alpha_j + \sum_{i=1}^t X_i, \beta_j + t - \sum_{i=1}^t X_i)$. Furthermore, a direct inspection of the definition of d_j^t shows that

$$d_j^t(P_j, \omega) = \begin{cases} b_0 & \text{if } P_j[X_0 | \mathcal{F}_t] < 1/2 \\ b_1 & \text{if } P_j[X_0 | \mathcal{F}_t] \geq 1/2 \end{cases} \quad (2.6)$$

To be precise, we notice that the case $P_j[X_0 | \mathcal{F}_t] = \frac{1}{2}$ cannot be decided upon according to (2.1). This problem could have been tackled via randomization, but this would have led to unnecessary cumbersome computations. The choice of selecting b_1 in this case is *ad hoc*, but creates no substantial harm.

In view of (2.4) and of the considerations above, we see that, for $j = 1, 2$,

$$P_j[X_0 | \mathcal{F}_t] = P_j[P_j[X_0 | \sigma(\Lambda)] | \mathcal{F}_t] = P_j[\Lambda | \mathcal{F}_t] = \frac{\alpha_j + \sum_{i=1}^t X_i}{\alpha_j + \beta_j + t}, \quad (2.7)$$

and therefore (2.6) can be written as

$$d_j^t(P_j) = \begin{cases} b_0 & \text{if } \sum_{i=1}^t X_i < (t + \beta_j - \alpha_j)/2, \\ b_1 & \text{otherwise} \end{cases}, \quad (2.8)$$

For the sake of clarity, we collect here the assumptions made so far:

- H1. The utilities of the shepherds and of the chief are given by (2.2) and (2.3), respectively.
- H2. The two shepherds maximize their respective expected utilities and the chief maximizes the social expected utility.
- H3. The chief decides in advance the number of times she is going to send out the scout before forcing the shepherds to act.
- H4. X_0, X_1, X_2, \dots is an exchangeable sequence, with respect to each of the measures P_0, P_1 and P_2 .
- H5. For all t , $P_j[X_0|\mathcal{F}_t]$ is an arithmetic mean of $P_j[X_0]$ and $\sum_{i=1}^t X_i/t$ as in (2.5).

3. The chief does not learn.

In this section, we consider the case in which the opinion of the chief does not evolve with time, i.e. is not affected by the information progressively reported by the scout. In other words, we assume

H6. $P_0^t = P_0, \forall t \in \mathbb{N}$.

Notice though that the chief knows that the opinions of the shepherds *do* evolve.

Since $P_0^t = P_0$, for all t , the infinite exchangeable sequence X_0, X_1, X_2, \dots is necessarily seen by the chief as an i.i.d. sequence of Bernoulli random variables with $P_0[X_1] = P_0(X_1 = 1) = 1 - P_0(X_1 = 0) = p$ ($p \in [0, 1]$).

Recall, by (2.8), that the j -th shepherd will choose path 0 at time t if $(\sum_{i=1}^t X_i/t) < 1/2 + (\beta_j - \alpha_j)/2t$. As t increases, the influence of the parameters (α_j, β_j) of the initial opinions decreases. By the law of large numbers (writing 1_A for the indicator function of the set A)

$$P_0 \left(\lim_{t \rightarrow \infty} \frac{\sum_{i=1}^t X_i}{t} < \frac{1}{2} \right) = \mathbf{1}_{(0, \frac{1}{2})}(p) \quad (3.1)$$

so that the chief expects the shepherds to choose, in the long run, b_0 if $p < \frac{1}{2}$ and b_1 otherwise.

In order to simplify the notation in the analysis of the expected utility, we write

$$\Psi(t) = P_0 [W(u_1, u_2, P_1, P_2, t, \cdot)] \quad (3.2)$$

and

$$\Psi(\infty) = \lim_{t \rightarrow \infty} \Psi(t) = P_0 \left[\lim_{t \rightarrow \infty} W(u_1, u_2, P_1, P_2, t, \cdot) \right]. \quad (3.3)$$

Notice that the interchange of limit and expectation is allowed, in view of the boundedness assumptions made.

The function $\Psi(t)$ represents the expected social utility if the chief forces the shepherds to act at time t . The following proposition quantifies its limit, when all available information is released.

Proposition 3.1. *Assume H1–H6, and let $\Psi(\infty)$ be as in (3.3). Then $\Psi(\infty) = \max(p, 1 - p)$.*

Proof. For $t \in \mathbb{N}$, let

$$\begin{aligned} A(t) &= P_0 \left(\sum_{i=1}^t X_i < \frac{t}{2} + \min_{j \in \{1,2\}} \left(\frac{\beta_j - \alpha_j}{2} \right) \right) \\ B(t) &= P_0 \left(\frac{t}{2} + \min_{j \in \{1,2\}} \left(\frac{\beta_j - \alpha_j}{2} \right) \leq \sum_{i=1}^t X_i < \frac{t}{2} + \max_{j \in \{1,2\}} \left(\frac{\beta_j - \alpha_j}{2} \right) \right) \\ C(t) &= P_0 \left(\sum_{i=1}^t X_i \geq \frac{t}{2} + \max_{j \in \{1,2\}} \left(\frac{\beta_j - \alpha_j}{2} \right) \right). \end{aligned} \quad (3.4)$$

It is easy to realize that $A(t)$, $B(t)$ and $C(t)$ represent the probabilities that, at time t , both shepherds choose path 0, that they go on different paths and that they both choose path 1, respectively. If they both select path 1, say, then the utility W is 1 with probability p and 0 with probability $1 - p$. If they choose different paths, then $W = \theta$ with probability one. Therefore

$$\Psi(t) = A(t)(1 - p) + B(t)\theta + C(t)p. \quad (3.5)$$

From (3.1) we deduce that

$$\begin{aligned} \lim_{t \rightarrow \infty} A(t) &= \mathbf{1}_{(0, \frac{1}{2})}(p) \\ \lim_{t \rightarrow \infty} B(t) &= 0 \\ \lim_{t \rightarrow \infty} C(t) &= \mathbf{1}_{[\frac{1}{2}, 1)}(p). \end{aligned} \quad (3.6)$$

The claim follows easily. ■

By the above proposition, the chief is sure that (in expectation) she can guarantee the society at least $\max(p, 1 - p)$. All that she has to do is just let

the information flow indefinitely. Stopping at some finite time will be rational only if, in this case, more than $\max(p, 1 - p)$ can be achieved. It is rather intuitive that the more unbalanced the opinion of the chief is (i.e. the further p is from $1/2$), the more she expects to be able to get for the society by letting information flow. This will have other implications, that will be made clear in Proposition 3.3.

We partition the interval $[0, 1]$ in two sets: \mathcal{I}_∞ , the set of possible values of p for which it is convenient to release as much information as possible, and its complement \mathcal{I}_f . More precisely,

$$\mathcal{I}_\infty = \mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, \theta) = \{p \in [0, 1] \mid \Psi(t) - \Psi(\infty) \leq 0, \forall t \in \mathbb{N}\} \quad (3.7)$$

and

$$\mathcal{I}_f = [0, 1] \setminus \mathcal{I}_\infty = \{p \in [0, 1] \mid \exists t < \infty \text{ such that } \Psi(t) > \Psi(\infty)\}. \quad (3.8)$$

In the following propositions, we shall study the structure of \mathcal{I}_∞ , and show that there are cases in which $\mathcal{I}_f \neq \emptyset$. In particular, in some instances it will be convenient not to release any information, namely $\Psi(0) \geq \Psi(t), \forall t \in \overline{\mathbb{N}}$, and in some other instances it will be optimal to release only a partial amount of information. First we determine a topological property of the set \mathcal{I}_∞ .

Proposition 3.2. *Assume H1-H6. The set \mathcal{I}_∞ defined in (3.7) is closed.*

Proof. We notice that Ψ depends on p through P_0 , and the function $p \mapsto \Psi(t) - \Psi(\infty)$ is continuous, since it depends continuously, through $A(t), B(t)$ and $C(t)$, on the cumulative distribution function of a binomial distribution with parameters t and p , which is continuous in p . Therefore, the set $\{p \in [0, 1] \mid \Psi(t) - \Psi(\infty) \leq 0\}$ is closed, and $\mathcal{I}_\infty = \bigcap_{t \in \mathbb{N}} \{p \in [0, 1] \mid \Psi(t) - \Psi(\infty) \leq 0\}$ is closed as well. ■

Next we show that if the chief has a sufficiently strong opinion about which of the two paths is the good one (p very close to 0 or 1), then it is optimal to release as much information as possible. How “strong” the opinion should be for this property to hold, depends on the degree of her risk-aversion.

Proposition 3.3. Assume H1–H6, and let \mathcal{I}_∞ be defined as in (3.7). Then, for all values of $\alpha_1, \beta_1, \alpha_2, \beta_2$ and θ ,

$$([0, 1 - \theta] \cup [\theta, 1]) \subset \mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, \theta).$$

Proof. It will be given in the Appendix.

Remark 3.4. If the social utility is convex, namely if $\theta \leq \frac{1}{2}$, then $\mathcal{I}_\infty = [0, 1]$. Therefore, only if the chief is risk-averse it might be optimal not to fully inform the shepherds. This is indeed the core of the whole issue: only the desire to avoid the risk of the extreme situation in which all the shepherds die might lead to withhold some beneficial additional information; and the degree of this desire is quantified by the concavity of the social utility.

In the following example we shall present a non-trivial case (i.e. $\theta > \frac{1}{2}$) in which $[0, 1 - \theta] \cup [\theta, 1] = \mathcal{I}_\infty$. We shall see in the sequel a case in which $[0, 1 - \theta] \cup [\theta, 1] \subsetneq \mathcal{I}_\infty$.

Example 3.5. Let $\theta > \frac{1}{2}$ and let $\alpha_1/(\alpha_1 + \beta_1) < \frac{1}{2} < \alpha_2/(\alpha_2 + \beta_2)$. Then $B(0) = 1$ and $\Psi(0) = \theta$. Thus

$$\Psi(0) - \Psi(\infty) = \theta - \max(p, 1 - p) = \min(\theta - p, \theta - 1 + p).$$

This is less than zero if and only if $p \geq \theta$ or $p \leq 1 - \theta$.

The next proposition shows that the learning processes of the shepherds are a crucial factor in deciding how much information ought to be released. If the opinions of the shepherds are too close and are not bound to lead to different actions, no matter what the X 's turn out to be, then it is optimal to release all the information. Conversely, if there is hope of diversification in the behavior of the shepherds ($B(t) > 0$), and if the chief is risk-averse ($\theta > 1/2$) and does not have extreme opinions (p close to $1/2$), then it might be optimal to withhold information.

Proposition 3.6. Assume H1–H6, and let \mathcal{I}_f and B be defined as in (3.8) and (3.4), respectively. Then $\mathcal{I}_f \neq \emptyset$ if and only if $\theta > 1/2$ and there exists a t_0 such that $B(t_0) > 0$. In such case, for some $\epsilon > 0$, $(1/2 - \epsilon, 1/2 + \epsilon) \subset \mathcal{I}_f$.

Proof. Let us prove first that, if $B(t) = 0 \forall t$, then $\mathcal{I}_f = \emptyset$. Let $p \geq \frac{1}{2}$. Then

$$\Psi(t) - \Psi(\infty) = A(t)(1 - p) + C(t)p - p = A(t)(1 - 2p) \leq 0 \forall t \in \mathbb{N}$$

so that $[\frac{1}{2}, 1] \subset \mathcal{I}_\infty$. Similarly, if $p < \frac{1}{2}$, then

$$\Psi(t) - \Psi(\infty) = A(t)(1 - p) + C(t)p - 1 + p = C(t)(2p - 1) < 0 \quad \forall t \in \mathbb{N}$$

so that also $[0, \frac{1}{2}] \subset \mathcal{I}_\infty$.

Conversely, let us show that, if $\exists t_0 : B(t_0) > 0$, then $\mathcal{I}_f \neq \emptyset$. Let $p = \frac{1}{2}, \theta > \frac{1}{2}$ and let $\alpha_1, \alpha_2, \beta_1, \beta_2$ be such that there exists t_0 with $B(t_0) > 0$. Then

$$\Psi(t_0) - \Psi(\infty) = \frac{1}{2}(A(t_0) + C(t_0)) + B(t_0)\theta - \frac{1}{2} = B(t_0)(\theta - \frac{1}{2}) > 0.$$

Therefore, $\frac{1}{2} \in \mathcal{I}_f(\alpha_1, \alpha_2, \beta_1, \beta_2, \theta)$. The claim follows from the fact that \mathcal{I}_∞ is closed. ■

Remark 3.7. $B(t) = 0 \forall t$ if, for instance, $\beta_1 - \alpha_1 \leq \beta_2 - \alpha_2$ and the interval $[\beta_1 - \alpha_1, \beta_2 - \alpha_2]$ contains no integers (positive or negative).

Remark 3.8. Although $[0, 1 - \theta] \cup [\theta, 1] \subset \mathcal{I}_\infty$, the set \mathcal{I}_∞ is not necessarily symmetric about $\frac{1}{2}$. It can be shown, for instance, that $0.6 \in \mathcal{I}_\infty(2, 4, 6, 12, \frac{7}{11})$, but $0.4 \in \mathcal{I}_f(2, 4, 6, 12, \frac{7}{11})$. This should not surprise, since the two agents are not identical. In the above case the initial opinion about the expected value of X_0 is the same for both shepherds: $P_1[X_0] = P_2[X_0] = 1/4$, but shepherd 1 is less opinionated than shepherd 2, namely the weight given to the empirical evidence in the updating is higher for the first shepherd than for the second one ($1 - \gamma_t^{(1)} = t/(2 + 6 + t)$, $1 - \gamma_t^{(2)} = t/(4 + 12 + t)$). This difference tends to vanish, as information accrues. If $p = 0.4$, then, at time 0 both shepherds will make the same decision that they would make in the long run (with P_0 -probability 1), i.e. they will go on path 0, guaranteeing an expected social utility of 0.6. So zero information in this case has the same effect as all information. But something more can be achieved, since, for some t , the probability that the two shepherds will make different decisions will be positive, but the probability that both will choose path 1 will be zero. Therefore, for these values of t , the expected social utility will be strictly more than 0.6 (and strictly less than $\theta = 7/11$). On the other hand, if $p = 0.6$, the initial decision (at time 0) of the shepherds will guarantee a social expected utility of 0.4, whereas the long run decision (after all information is released) will guarantee 0.6. At some

time t the event that the two shepherds make different decisions will have positive probability, but the event that they will both go on path 0 will always have positive probability. For these particular values of the parameters, it will always be safer for the chief to let information flow and get 0.6, rather than betting on the event of different decisions of the shepherds and risking instead that they both go on path 0.

Remark 3.9. If $B(0) = 0$ and if $B(t) > 0$ for some $t \in \mathbb{N}$, then the Proposition above implies that $\frac{1}{2} \in \mathcal{I}_I$ and, for $p = \frac{1}{2}$, there exists an optimal strictly positive finite t_0 , i.e. it is optimal to release *some* information (as opposed to *none* or *all*). Furthermore, notice that whenever $B(0) \neq 0$, then $B(0) = 1$, and this happens if

$$\min_{j \in \{1,2\}} (\beta_j - \alpha_j) < 0 \leq \max_{j \in \{1,2\}} (\beta_j - \alpha_j).$$

Example 3.10. This clarifies the content of the preceding remark. Let $p = \frac{1}{2}, \theta > \frac{1}{2}, \alpha_1 = 2, \beta_1 = 8, \alpha_2 = 4, \beta_2 = 6$. It will be shown in the Appendix that $\Psi(2k) > \Psi(2k + 1), \forall k \in \mathbb{N}$, and that

$$\frac{\Psi(2k)}{\Psi(2k + 2)} \begin{cases} < 1 & k = 0, 1, 2 \\ = 1 & k = 3 \\ > 1 & k > 3. \end{cases} \quad (3.9)$$

Therefore, it is optimal to release either 6 or 8 units of information. Notice, incidentally, that $\Psi(0) = \Psi(\infty) = 0.5$ and $\Psi(6) = \Psi(8) \simeq 0.574574$.

4. The chief learns.

Now the flow of information will change not only the opinion of the shepherds, but that of the chief as well. She updates her opinion like the shepherds, namely she considers X_0, X_1, X_2, \dots exchangeable and such that

$$\text{H6bis.} \quad P_0[X_0 \mid \mathcal{F}_t] = \frac{\alpha_0 + \sum_{i=1}^t X_i}{\alpha_0 + \beta_0 + t} \quad \alpha_0, \beta_0 > 0.$$

This is equivalent to assuming X_0, X_1, X_2, \dots P_0 -i.i.d. conditionally on a parameter Λ , which is distributed according to a Beta(α_0, β_0). Nonetheless, the chief is still forced to decide at the beginning the amount of information she is going to release. Let

$$\begin{aligned} A_\Lambda(t) &= P_0 \left(\sum_{i=1}^t X_i < \frac{t}{2} + \min_{j \in \{1,2\}} \left(\frac{\beta_j - \alpha_j}{2} \right) \mid \sigma(\Lambda) \right) \\ B_\Lambda(t) &= P_0 \left(\frac{t}{2} + \min_{j \in \{1,2\}} \left(\frac{\beta_j - \alpha_j}{2} \right) \leq \sum_{i=1}^t X_i < \frac{t}{2} + \max_{j \in \{1,2\}} \left(\frac{\beta_j - \alpha_j}{2} \right) \mid \sigma(\Lambda) \right) \\ C_\Lambda(t) &= P_0 \left(\sum_{i=1}^t X_i \geq \frac{t}{2} + \max_{j \in \{1,2\}} \left(\frac{\beta_j - \alpha_j}{2} \right) \mid \sigma(\Lambda) \right). \end{aligned}$$

See formula (3.4) for related definitions. Let also $p_0 = P_0[X_0] = \alpha_0/(\alpha_0 + \beta_0)$. Then (3.5) becomes, in this setup,

$$\begin{aligned} \Psi(t) &= P_0 \left[P_0 [W(u_1, u_2, P_1, P_2, t, \cdot) \mid \Lambda] \right] \\ &= P_0 [A_\Lambda(t)(1 - \Lambda) + B_\Lambda(t)\theta + C_\Lambda(t)\Lambda] \end{aligned} \quad (4.1)$$

Hence

$$\Psi(0) = \begin{cases} P_0[1 - \Lambda] = 1 - p_0 & \text{if } \min_{j \in \{1,2\}} (\beta_j - \alpha_j) > 0 \\ \theta & \text{if } \min_{j \in \{1,2\}} (\beta_j - \alpha_j) \leq 0 < \max_{j \in \{1,2\}} (\beta_j - \alpha_j) \\ P_0[\Lambda] = p_0 & \text{if } \max_{j \in \{1,2\}} (\beta_j - \alpha_j) \leq 0. \end{cases}$$

By letting $t \rightarrow \infty$, we obtain the following result about the limit of the expected social utility. Let $B(\alpha, \beta) = \Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha + \beta)$ be the beta integral, and let

$$B_x(\alpha, \beta) = \int_0^x \lambda^{\alpha-1} (1 - \lambda)^{\beta-1} d\lambda$$

be the incomplete beta function. Let also

$$I_x(\alpha, \beta) = \frac{B_x(\alpha, \beta)}{B(\alpha, \beta)}$$

be the cumulative distribution function of a Beta distribution with parameters α and β , evaluated at x .

Proposition 4.1. Assume H1–H5 and H6bis. Let Ψ be as in (4.1). Then

$$\Psi(\infty) = \frac{\beta_0}{\alpha_0 + \beta_0} I_{\frac{1}{2}}(\alpha_0, \beta_0 + 1) + \frac{\alpha_0}{\alpha_0 + \beta_0} \left[1 - I_{\frac{1}{2}}(\alpha_0 + 1, \beta_0) \right] \quad (4.2).$$

Proof. By conditioning on Λ , we can use (3.6) and we obtain

$$\begin{aligned} \Psi(\infty) &= P_0 \left[(1 - \Lambda) \mathbf{1}_{[0, \frac{1}{2})}(\Lambda) + \Lambda \mathbf{1}_{[\frac{1}{2}, 1]}(\Lambda) \right] \quad (4.3) \\ &= \int_0^{\frac{1}{2}} \frac{\lambda^{\alpha_0 - 1} (1 - \lambda)^{\beta_0}}{B(\alpha_0, \beta_0)} d\lambda + \int_{\frac{1}{2}}^1 \frac{\lambda^{\alpha_0} (1 - \lambda)^{\beta_0 - 1}}{B(\alpha_0, \beta_0)} d\lambda \\ &= \frac{B_{\frac{1}{2}}(\alpha_0, \beta_0 + 1)}{B(\alpha_0, \beta_0)} + \frac{B(\alpha_0 + 1, \beta_0) - B_{\frac{1}{2}}(\alpha_0 + 1, \beta_0)}{B(\alpha_0, \beta_0)} \\ &= \frac{B(\alpha_0, \beta_0 + 1)}{B(\alpha_0, \beta_0)} I_{\frac{1}{2}}(\alpha_0, \beta_0 + 1) + \frac{B(\alpha_0 + 1, \beta_0)}{B(\alpha_0, \beta_0)} - \frac{B(\alpha_0 + 1, \beta_0)}{B(\alpha_0, \beta_0)} I_{\frac{1}{2}}(\alpha_0 + 1, \beta_0) \end{aligned}$$

and the claim follows easily. ■

Remark 4.2. The case of Section 3 can be obtained by letting $\alpha_0 \rightarrow \infty, \beta_0 \rightarrow \infty$, with $\alpha_0/(\alpha_0 + \beta_0) = p_0$. In fact, both $I_{\frac{1}{2}}(\alpha_0, \beta_0 + 1)$ and $I_{\frac{1}{2}}(\alpha_0 + 1, \beta_0)$ converge to $\mathbf{1}_{[0, \frac{1}{2})}(p_0)$, which yields the value at $x = \frac{1}{2}$ of the cumulative distribution function of a random variable degenerate at p_0 . Therefore, $\Psi(\infty) \rightarrow \max(p_0, 1 - p_0)$.

Remark 4.3. As $\alpha_0, \beta_0 \rightarrow 0$ (with $\alpha_0/(\alpha_0 + \beta_0) = p_0$) the law of the sequence $\{X_n \mid n \in \mathbb{N}\}$ tends to the maximum of its Fréchet class, namely for every $n \in \mathbb{N}$ we have, in the limit,

$$P_0(X_0 = X_1 = X_2 = \dots = X_n = 1) = p_0 = 1 - P_0(X_0 = X_1 = X_2 = \dots = X_n = 0)$$

This can be easily shown considering that, with respect to P_0 ,

$$\begin{aligned} \rho(X_1, X_2) &= \frac{\text{Cov}(X_1, X_2)}{\text{Var}(X_1)} = \frac{P_0[\text{Cov}(X_1, X_2 \mid \Lambda)] + \text{Cov}(P_0[X_1 \mid \Lambda], P_0[X_2 \mid \Lambda])}{P_0[\text{Var}(X_1 \mid \Lambda)] + \text{Var}(P_0[X_1 \mid \Lambda])} \\ &= \frac{\text{Var}(\Lambda)}{P_0[\Lambda(1 - \Lambda)] + \text{Var}(\Lambda)} = \frac{1}{\alpha_0 + \beta_0 + 1} \end{aligned}$$

so that $\lim_{\alpha_0, \beta_0 \rightarrow 0} \rho(X_1, X_2) = 1$.

In this case all the random variables of the sequence are equal P_0 -almost surely, and $\Psi(\infty) = 1$, since the chief expects to see (almost surely) a sequence of identical

random variables, all equal to X_0 . Therefore she knows that, whatever the initial opinion of the shepherds is (provided it is nondegenerate), their distribution for X_0 will converge to a degenerate distribution on the correct value. This will render the expected social utility equal to its maximum possible value.

Remark 4.4. We can show that the learning chief can guarantee in the long run a higher social utility than the nonlearning one, i.e. $\Psi(\infty) \geq \max(p_0, 1 - p_0)$, with $\Psi(\infty)$ given in Proposition 4.1. In fact, by formula (4.3) and Jensen inequality

$$\begin{aligned} \Psi(\infty) &= P_0[(1 - \Lambda)\mathbf{1}_{[0, \frac{1}{2})}(\Lambda)] + P_0[\Lambda\mathbf{1}_{[\frac{1}{2}, 1]}(\Lambda)] \\ &= P_0[\max(\Lambda, (1 - \Lambda))] \geq \max(P_0[1 - \Lambda], P_0[\Lambda]) = \max(p_0, 1 - p_0) \end{aligned} \quad (4.4)$$

with strict inequality holding whenever Λ is nondegenerate. The above phenomenon is not due to the fact that a learning chief can take advantage of the information that becomes available. In fact her decision is made once for all at the beginning and she cannot change her mind subsequently, no matter what she sees. The real reason is the following. A nonlearning chief knows that $\Lambda = p_0$ a.s., whereas a learning chief only knows that $P_0[\Lambda] = p_0$ and therefore Λ can assume different values, smaller and larger than p_0 . The decision of the learning chief will therefore involve an integration of the social utility over the possible values of Λ . Since the social utility is a convex function of Λ , its expected value will be larger than the function computed at the expected value of Λ (p_0). Similar phenomena in the framework of exchangeable coin tossings have been studied by Scarsini (1984).

We shall show now that whenever the social utility function is convex, it is more convenient to release as much information as possible.

Proposition 4.5. *Assume H1–H5 and H6bis. Let Ψ be as in (4.1), and let $\theta \leq \frac{1}{2}$. Then*

$$\Psi(t) \leq \Psi(\infty) \quad \forall \alpha_0, \beta_0 \in \mathbb{R}_+, \forall t \in \mathbb{N}.$$

Proof. It will be given in the Appendix.

Remark 4.6. Arguments similar to those used to prove Proposition 4.5 lead to a

lower bound for Ψ when the social utility is concave. For $\theta \geq \frac{1}{2}$,

$$\begin{aligned} \Psi(t) &\geq \frac{1}{B(\alpha_0, \beta_0)} \left\{ \int_0^{\frac{1}{2}} \lambda^{\alpha_0} (1-\lambda)^{\beta_0-1} d\lambda + \int_{\frac{1}{2}}^1 \lambda^{\alpha_0-1} (1-\lambda)^{\beta_0} d\lambda \right\} \\ &= \frac{\alpha_0}{\alpha_0 + \beta_0} I_{\frac{1}{2}}(\alpha_0 + 1, \beta_0) + \frac{\beta_0}{\alpha_0 + \beta_0} [1 - I_{\frac{1}{2}}(\alpha_0, \beta_0 + 1)]. \end{aligned} \quad (4.5)$$

Notice that, as $\alpha_0, \beta_0 \rightarrow \infty$, with $\alpha_0/(\alpha_0 + \beta_0) = p_0$, then the lower bound becomes $\min(p_0, 1 - p_0)$.

Example 4.7. As in the previous section, we exhibit a case in which it is optimal not to release all available information. Let $\alpha_0 = \beta_0 = 5, \theta = \frac{3}{4}, \beta_1 - \alpha_1 = 4, \beta_2 - \alpha_2 = 8$. It can be shown that $\Psi(5) - \Psi(\infty) = 0.0088133$. The existence of such an example is not at all surprising, given Example 3.10 for the nonlearning chief and Remark 4.6: The behavior of the nonlearning chief can be obtained as a limit of the behavior of a learning chief. Some continuity arguments provide the result.

5. Conclusions.

We have considered a situation in which a central planner (the chief) aims at influencing the decisions made by several agents (the shepherds), each of which maximizes his own expected utility with respect to his own subjective probability. The planner maximizes the expected social utility function, which is increasing and symmetric in the utilities of the various agents. Typically the social utility function will be concave, so some diversification in the strategies of the agents will be socially beneficial. Several scenarios could be conceived. One of them is the dictatorial situation where the planner can select the strategy for each agent. According to the concavity of the social utility function, the strategies that the planner chooses for the agents will be more or less diversified. If the dictator is benevolent, she will assign the strategies taking into account the agents' utility functions and subjective probabilities, so that, even if some agents will not be assigned their first choice strategy, still they will not be completely unhappy.

In a different scenario the agents cannot be forced to act against their self interest, but the chief exerts her influence on the agents indirectly by correlating

information. In this scenario, different possibilities can be examined. For instance the chief could provide the agents with different streams of information. If the information is censored, namely the chief releases a random bit of information to a certain shepherd only if the value of that bit is the desired one, then we are *de facto* back to the trivial case of dictatorship. Even if the information is not censored, still the chief could decide to give different agents different amounts of information (or the same amount of bits, but from different sets of data). This would allow the chief more freedom than the case we have considered in this paper, where the information that is released is not only uncensored, but also the same for all the agents. Surprisingly, even in this very restrictive scenario for the planner, the flow of information can be used to maximize the social expected utility.

Spreading information has the effect of making the probability distributions of the agents more concentrated on the same states of nature. This is due to the fact that the information is the same for all agents and the conditions for merging of opinions are satisfied. Information is free, so it would seem that the more the better. This is actually the case when the accruing of information leads to the complete elimination of uncertainty. When it doesn't (as in the situation we examined), the merging effect and the effect of improving the opinion of each agent conflict. As information increases, the random variable that determines the social utility (the number of agents who make the right decision) tend to have higher mean, but also to be more dispersed over extreme values: if all the agents make the same decision, they are either all right or all wrong. Therefore, if the utility is very concave, then the uniformity of the agents' decision could become socially dangerous, even if on the average the number of agents that get it right increases.

We have considered a case where the opinion of the chief is affected by observations of the information released to the agents, and a case where it is not. In both cases we have shown the existence of situations where it is socially better to withdraw some information from the public. Nevertheless, even if the planner is required to make her decision once and for all at the beginning of the game, the case of the learning chief and the case of the nonlearning chief differ somehow. When the chief learns from experience, she can achieve a higher expected social

utility than when she does not learn. This is not due to the learning itself, since the decision is made before the information is obtained, but rather to the fact that the learning chief, being more uncertain, sees the observations as positively correlated and therefore assigns higher probability to the event that all the agents make the same decision and hence is less hopeful in diversification.

As a final comment it is worth noticing that the assumption of concavity of the social utility function is crucial for the optimality of withdrawing information. We have proved that if the social utility function is not concave, then it is optimal to release full information, since in this case diversification of the agents' strategies would not be of any use. It's only when diversification is needed that withdrawing information could be useful as an instrument to achieve it.

Acknowledgments. We thank Ido Erev for inspiring discussions, Isaco Meilijson for helpful suggestions and Giovanni Cespa for a careful reading of a first draft. We are especially indebted to an associate editor for many useful comments.

Appendix.

Notation. In the sequel, we shall write $\Psi_{p,\theta}$ instead of Ψ whenever it is appropriate to emphasize the dependence on p and θ .

Proof of Proposition 3.3. The proof relies on the following lemmata.

Lemma A.1. *If $\theta_1 < \theta_0$, then*

$$\mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, \theta_0) \subset \mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, \theta_1).$$

Proof. By (3.5) and Proposition 3.1, we see that the function $\theta \mapsto \Psi_{p,\theta}(t) - \Psi_{p,\theta}(\infty)$ is increasing. If, furthermore, $p \in \mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, \theta_0)$, then

$$\Psi_{p,\theta_1}(t) \leq \Psi_{p,\theta_0}(t) \leq \Psi_{p,\theta_0}(\infty) = \Psi_{p,\theta_1}(\infty) \quad \forall t \in \mathbb{N},$$

so that $p \in \mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, \theta_1)$. ■

Lemma A.2. *For all values of $\alpha_1, \beta_1, \alpha_2, \beta_2$ and θ ,*

$$\theta \in \mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, \theta) \tag{A.1}$$

and

$$1 - \theta \in \mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, \theta). \quad (\text{A.2})$$

Proof. Let us start with the case $\theta \geq \frac{1}{2}$. If $p = \theta$, then $\Psi_{p,\theta}(\infty) = \theta$. Recalling that $A(t) + B(t) + C(t) = 1$, we get

$$\begin{aligned} \Psi_{p,\theta}(t) - \Psi_{p,\theta}(\infty) &= A(t)(1 - \theta) + B(t)\theta + C(t)\theta - \theta \\ &= \theta(B(t) + C(t) - A(t) - 1) + A(t) \\ &= -2\theta A(t) + A(t) \leq 0 \quad \forall t \in \mathbb{N} \end{aligned}$$

which implies (A.1).

If $p = 1 - \theta$, then again $\Psi_{p,\theta}(\infty) = \theta$, and

$$\begin{aligned} \Psi_{p,\theta}(t) - \Psi_{p,\theta} &= A(t)\theta + B(t)\theta + C(t)(1 - \theta) - \theta \\ &= C(t)(1 - 2\theta) \leq 0 \quad \forall t \in \mathbb{N}. \end{aligned}$$

Thus, (A.2) is also proved.

In the case $\theta < \frac{1}{2}$, we have $\Psi_{1-\theta,\theta}(\infty) = \Psi_{\theta,\theta}(\infty) = 1 - \theta$, and the proof is similar. ■

Proof of Proposition 3.3. Let $p > \theta$. By Lemma A.2, $p \in \mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, p)$. Thus, we have proved that $[\theta, 1] \subset \mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, \theta)$.

It remains to prove that $p < 1 - \theta \Rightarrow p \in \mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, \theta)$. The arguments are similar to those above: by Lemma A.2, $p \in \mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, 1 - p)$. Since $\theta < 1 - p$, Lemma A.1 implies $p \in \mathcal{I}_\infty(\alpha_1, \beta_1, \alpha_2, \beta_2, \theta)$. ■

Example 3.10. Recalling (2.7) and (2.8),

$$d_1^t(P_1) = \begin{cases} b_0 & \text{if } \sum_{i=1}^t X_i < \frac{t}{2} + 3 \\ b_1 & \text{otherwise} \end{cases}$$

and

$$d_2^t(P_2) = \begin{cases} b_0 & \text{if } \sum_{i=1}^t X_i < \frac{t}{2} + 1 \\ b_1 & \text{otherwise.} \end{cases}$$

Thus

$$u_0(u_1, u_2, d_1^t(P_1), d_2^t(P_2)) = \begin{cases} 0 & \text{if } \sum_{i=1}^t X_i < \frac{t}{2} + 1 \text{ and } X_0 = 1 \\ & \text{or } \sum_{i=1}^t X_i \geq \frac{t}{2} + 3 \text{ and } X_0 = 0 \\ 1 & \text{if } \sum_{i=1}^t X_i < \frac{t}{2} + 1 \text{ and } X_0 = 0 \\ & \text{or } \sum_{i=1}^t X_i \geq \frac{t}{2} + 3 \text{ and } X_0 = 1 \\ \theta & \text{otherwise.} \end{cases}$$

Therefore

$$\begin{aligned} \Psi_{\frac{1}{2}, \theta}(t) &= 1 \cdot \frac{1}{2} \left\{ P_0 \left(\sum_{i=1}^t X_i < \frac{t}{2} + 1 \right) + P_0 \left(\sum_{i=1}^t X_i \geq \frac{t}{2} + 3 \right) \right\} \\ &\quad + \theta \left\{ 1 - \frac{1}{2} P_0 \left(\sum_{i=1}^t X_i < \frac{t}{2} + 1 \right) - \frac{1}{2} P_0 \left(\sum_{i=1}^t X_i \geq \frac{t}{2} + 3 \right) \right. \\ &\quad \left. - \frac{1}{2} P_0 \left(\sum_{i=1}^t X_i < \frac{t}{2} + 1 \right) - \frac{1}{2} P_0 \left(\sum_{i=1}^t X_i \geq \frac{t}{2} + 3 \right) \right\} \\ &= \theta + \left[\frac{1}{2} - \theta \right] \left\{ P_0 \left(\sum_{i=1}^t X_i < \frac{t}{2} + 1 \right) + P_0 \left(\sum_{i=1}^t X_i \geq \frac{t}{2} + 3 \right) \right\}. \end{aligned}$$

Denoting by $\lceil x \rceil$ the ceiling of x (i.e. the smallest integer greater or equal to x), we may write

$$\Psi_{\frac{1}{2}, \theta}(t) = \theta + \left(\frac{1}{2} - \theta \right) \left\{ \sum_{j=0}^{\lceil \frac{t}{2} \rceil} \frac{1}{2^t} \binom{t}{j} + \sum_{j=\lceil \frac{t}{2} \rceil + 3}^t \frac{1}{2^t} \binom{t}{j} \right\}.$$

Let us distinguish the cases of t even and t odd:

$$\begin{aligned} \Psi_{\frac{1}{2}, \theta}(2k) &= \theta + \left(\frac{1}{2} - \theta \right) \frac{1}{2^{2k}} \left\{ \sum_{j=0}^k \binom{2k}{j} + \sum_{j=k+3}^{2k} \binom{2k}{j} \right\} \\ &= \theta + \left(\frac{1}{2} - \theta \right) \left\{ 1 - \frac{1}{2^{2k}} \binom{2k+1}{k+2} \right\} \\ \Psi_{\frac{1}{2}, \theta}(2k+1) &= \theta + \left(\frac{1}{2} - \theta \right) \frac{1}{2^{2k+1}} \left\{ \sum_{j=0}^{k+1} \binom{2k+1}{j} + \sum_{j=k+4}^{2k+1} \binom{2k+1}{j} \right\} \\ &= \theta + \left(\frac{1}{2} - \theta \right) \left\{ 1 - \frac{1}{2^{2k+1}} \binom{2k+2}{k+3} \right\}. \end{aligned}$$

It is now easy to verify that $\Psi_{\frac{1}{2}, \theta}(2k) > \Psi_{\frac{1}{2}, \theta}(2k+1)$, $\forall k \in \mathbb{N}$ and that (3.9) holds.

Proof of Proposition 4.5. Define

$$z = \left\lceil \frac{t}{2} + \min_{j \in \{1,2\}} \left(\frac{\beta_j - \alpha_j}{2} \right) \right\rceil \quad \text{and} \quad w = \left\lceil \frac{t}{2} + \max_{j \in \{1,2\}} \left(\frac{\beta_j - \alpha_j}{2} \right) \right\rceil.$$

By (4.1), we have

$$\begin{aligned} \Psi(t) = \int_0^1 & \left\{ \sum_{j=0}^{z-1} \binom{t}{j} \lambda^j (1-\lambda)^{t-j+1} + \theta \sum_{j=z}^{w-1} \binom{t}{j} \lambda^j (1-\lambda)^{t-j} \right. \\ & \left. + \sum_{j=w}^t \binom{t}{j} \lambda^{j+1} (1-\lambda)^{t-j} \right\} \frac{\lambda^{\alpha_0-1} (1-\lambda)^{\beta_0-1}}{B(\alpha_0, \beta_0)} d\lambda \end{aligned} \quad (\text{A.3}).$$

By using the well-known formula for the cumulative distribution function of the binomial distribution

$$\sum_{j=0}^{z-1} \binom{n}{j} p^j (1-p)^{n-j} = 1 - I_p(z, n - z + 1),$$

(see, for example, Feller (1970) or Selby (1967)), formula (A.3) reduces to

$$\begin{aligned} \Psi(t) = \int_0^1 & \left\{ (1-\lambda)[1 - I_\lambda(z, t - z + 1)] + \theta [I_\lambda(z, t - z + 1) - I_\lambda(w, t - w + 1)] \right. \\ & \left. + \lambda I_\lambda(w, t - w + 1) \right\} \frac{\lambda^{\alpha_0-1} (1-\lambda)^{\beta_0-1}}{B(\alpha_0, \beta_0)} d\lambda. \end{aligned} \quad (\text{A.4})$$

When $\theta = \frac{1}{2}$, we can bound θ from above with $1 - \lambda$ for $\lambda < \frac{1}{2}$ and with λ for $\lambda \geq \frac{1}{2}$. Thus

$$\begin{aligned} \Psi(t) \leq & \int_0^{\frac{1}{2}} \left\{ (1-\lambda)[1 - I_\lambda(w, t - w + 1)] + \lambda I_\lambda(w, t - w + 1) \right\} \frac{\lambda^{\alpha_0-1} (1-\lambda)^{\beta_0-1}}{B(\alpha_0, \beta_0)} d\lambda \\ & + \int_{\frac{1}{2}}^1 \left\{ (1-\lambda)[1 - I_\lambda(z, t - z + 1)] + \lambda I_\lambda(z, t - z + 1) \right\} \frac{\lambda^{\alpha_0-1} (1-\lambda)^{\beta_0-1}}{B(\alpha_0, \beta_0)} d\lambda. \end{aligned}$$

Since $x \mapsto (1-\lambda)(1-x) + \lambda x$ is decreasing for $\lambda \in [0, \frac{1}{2})$ and increasing for $\lambda \in (\frac{1}{2}, 1]$, we can bound the quantities between braces with their maxima, namely $1 - \lambda$ in the first integral and λ in the second one. Therefore

$$\Psi(t) \leq \int_0^{\frac{1}{2}} \frac{\lambda^{\alpha_0-1} (1-\lambda)^{\beta_0}}{B(\alpha_0, \beta_0)} d\lambda + \int_{\frac{1}{2}}^1 \frac{\lambda^{\alpha_0} (1-\lambda)^{\beta_0-1}}{B(\alpha_0, \beta_0)} d\lambda = \Psi(\infty). \quad (\text{A.5})$$

Hence, the claim is proved for $\theta = \frac{1}{2}$. The case $\theta < \frac{1}{2}$ can be dealt with by considering that Lemma A.1 applies also to the situation in which the chief learns: The function $\theta \mapsto \Psi_{p,\theta}(t) - \Psi_{p,\theta}(\infty)$ is increasing also for the learning chief (see formulas (A.4) and (4.2)). Therefore, if $\theta_1 < \theta_0$ and if α_0 and β_0 are such that $\Psi_{p,\theta_0}(t) \leq \Psi_{p,\theta_0}(\infty) \forall t \in \mathbb{N}$, then $\Psi_{p,\theta_1}(t) \leq \Psi_{p,\theta_1}(\infty)$, too. ■

References.

- Blackwell, D. and Dubins, L. (1962) Merging of opinions with increasing information, *Annals of Mathematical Statistics*, **38**, 882–886.
- de Finetti, B. (1930) Funzione caratteristica di un fenomeno aleatorio, *Memorie della Reale Accademia Nazionale dei Lincei*, **4**, 86–133.
- de Finetti, B. (1937) La prévision, ses lois logiques, ses sources subjectives, *Annales de l'Institut Henri Poincaré*, **7**, 1–68.
- Diaconis, P. and Ylvisaker, D. (1979) Conjugate priors for exponential families, *Annals of Statistics*, **7**, 269–281.
- Erev, I., Wallsten, T.S. and Neal, M.M (1991) Vagueness, ambiguity, and the cost of mutual understanding, *Psychological Science*, **2**, 321–324.
- Feller, W. (1970) *Introduction to Probability Theory and its Applications*, third edition, Wiley, New York.
- Kadane, J.B. (1993) Several Bayesians: A review, *Test*, **2**, 1–32 (with discussion).
- Kalai, E. and Lehrer, E. (1991) Weak and strong merging of opinions, *Journal of Mathematical Economics*, to appear.
- Lehrer, E. and Smorodinsky, R. (1993) Compatible measures and merging, School of Mathematical Sciences, Tel Aviv University.
- Scarsini, M. (1988) On a simple sequential decision problem, *Statistica Neerlandica*, **42**, 29–35.
- Selby, S.M. (1967) *Standard Mathematical Tables*, 15th edition, The Chemical Rubber, Cleveland.