Feasible Joint Posterior Beliefs (Through Examples)

ITAI ARIELI
Technion - Israel Institute of Technology
and
YAKOV BABICHENKO
Technion - Israel Institute of Technology
and
FEDOR SANDOMIRSKIY
Technion - Israel Institute of Technology; Higher School of Economics St. Petersburg
and
OMER TAMUZ
Caltech - California Institute of Technology

Through a sequence of examples, we survey the main results of "Feasible Joint Posterior Beliefs" [Arieli, Babichenko, Sandomirskiy, Tamuz 2021]. A group of agents share a common prior distribution regarding a binary state, and observe some information structure. What are the possible joint distributions of their posteriors? We discuss feasibility of product distributions, correlation of posteriors in feasible distributions, extreme feasible distributions and the characterization of feasibility in terms of a "no-trade" condition.

Categories and Subject Descriptors: J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*; G.3 [**Mathematics of Computing**]: Probability and Statistics; I.2.11 [**Computing Methodologies**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms: Economics
Additional Key Words and Phrases: information design, Bayesian persuasion, social learning, information structure, posterior belief, splitting lemma, no-trade theorem, extreme point, joint distribution

Consider a single agent, Alice, who is interested in an uncertain state of the world $\omega \in \{0, 1\}$, for which she has a prior probability $p = \mathbb{P}(\omega = 1)$. The state ω may indicate whether a piece of art being auctioned is a forgery, which of the two candidates will win an election, or whether a company will default this year. Alice receives some information about ω, in the form of a signal s_A taking values in a set S_A. For example, a signal can be a binary message, a real number, or it may contain information obtained by Alice during a dynamic learning process, e.g., a trajectory of stock market quotes. The signal is drawn at random from one of two distribution, depending on the state: π_0 when $\omega = 0$, and π_1 when $\omega = 1$. Alice is Bayes-rational and, hence, upon receiving the signal, updates her belief according

Authors' addresses: iarieli@technion.ac.il, yakovbab@tx.technion.ac.il, fedor.sandomirskiy@gmail.com, tamuz@caltech.edu

to Bayes' Law: her posterior belief about the high state $\omega = 1$ is[1]

$$p_A(s_A) = \frac{p \cdot \pi_1(s_A)}{p \cdot \pi_1(s_A) + (1-p) \cdot \pi_0(s_A)}.$$

The tuple (S_A, p, π_0, π_1) is referred to as an *information structure*. Unconditional on the state, the distribution of the signal is $\pi = p\pi_1 + (1-p)\pi_0$. Since Alice's posterior is a function of the signal, it is also random; its distribution is determined by π, through the application of Bayes' Law, which maps the signal to the posterior. We denote this (unconditional) distribution, a probability measure on the interval $[0, 1]$, by P. And we say that P is *implemented by* the information structure (S_A, p, π_0, π_1).

Which posterior belief distributions P are feasible, i.e., implementable by some information structure? The prior p is assumed to be fixed. In this single agent case, feasibility is determined by the so-called *martingale condition*: the expectation of Alice's posterior must be equal the prior, $\mathbb{E}_P[p_A] = p$. This fundamental result is known as the Splitting Lemma [Blackwell 1951; Aumann and Maschler 1995] and is a key tool in the theories of Bayesian persuasion [Kamenica and Gentzkow 2011] and of games with incomplete information [Aumann and Maschler 1995].

In [Arieli et al. 2020] and [Arieli et al. 2021], we study the question of feasibility for more than one agent. In this letter we focus on the case of two agents, Alice and Bob, who share a common prior p and use possibly different sources of information about the state ω. This is modelled by private signals, $s_A \in S_A$ for Alice and $s_B \in S_B$ for Bob, jointly distributed according to either π_0 or π_1, depending on the state. These distributions are now probability measures on the product $S_A \times S_B$. Both agents calculate their posteriors as in the single-agent case, each using her or his signal. We now denote by P the joint distribution of their posteriors. This is a distribution on the unit square $[0, 1] \times [0, 1]$, which is implemented by the information structure $(S_A, S_B, p, \pi_0, \pi_1)$.

The main question that we tackle is the inverse question: given a distribution P on the unit square, does there exist an information structure that implements P? If so, we say that P is feasible. For simplicity, below we focus on the symmetric prior $p = 1/2$ and, by feasibility, mean feasibility for this prior.

Example 1. Denote by $\delta_{(\cdot)}$ the Dirac measure (a point mass) and consider the distribution $P = \frac{1}{2}\delta_{(0,0)} + \frac{1}{2}\delta_{(1,1)}$, placing equal weight on the two main-diagonal corners. This distribution represents a situation where both Alice and Bob are either sure that the state is $\omega = 1$ (their posteriors are $p_A = p_B = 1$) or that $\omega = 0$ ($p_A = p_B = 0$). This distribution is feasible as it can be generated by revealing the state to both Alice and Bob. Formally, it is implemented by the information structure with $S_A = S_B = \{0, 1\}$, $\pi_0(0, 0) = 1$ and $\pi_1(1, 1) = 1$.

The martingale condition on the marginals, $\mathbb{E}_P[p_A] = \mathbb{E}_P[p_B] = p$, is obviously a necessary condition for feasibility, but it is insufficient as demonstrated by the following example.

[1] For a continuum of possible signals and non-atomic π_0, π_1, the ratio is not well-defined and has to be replaced by the Radon–Nikodym derivative.

Example 2. The distribution $P = \frac{1}{2}\delta_{(0,1)} + \frac{1}{2}\delta_{(1,0)}$, placing equal weight on the anti-diagonal corners of the square, satisfies the martingale condition for the prior $p = \frac{1}{2}$. However, P is infeasible. This distribution corresponds to a situation, where one agent is certain that the state is $\omega = 1$, while the other one is sure that $\omega = 0$. Such combination of beliefs cannot happen with positive probability. Indeed, from the definition of the posteriors, it is easy to deduce that $\mathbb{P}(\omega = 1 \mid p_A = x) = x$ for almost all $x \in [0,1]$, i.e., conditional on Alice's belief being x, the probability of the high state also equals x. Hence, conditionally on $p_A = 1$, the state $\omega = 1$ almost surely and, similarly, $\omega = 0$ conditionally on $p_B = 0$. Thus the intersection of events $\{p_A = 1\}$ and $\{p_B = 0\}$ cannot have positive probability, which proves the infeasibility of P.

Feasible distributions defying intuition

Example 3. Is the uniform distribution on $[0,1]^2$ feasible? At first glance, one might expect a negative answer. For this distribution, the posterior p_A of Alice reveals to her no information about the posterior p_B of Bob (conditionally on p_A, the distribution of p_B remains uniform) and vice versa. Intuitively, since Alice and Bob form beliefs about the same underlying state, it would be natural to expect positive correlation between posteriors: when p_A is high, the state $\omega = 1$ is more likely and, hence, p_B is also more likely to be high. This intuition turns out to be wrong.

[Gutmann et al. 1991, Example 2] describe the following information structure implementing the uniform distribution. The sets of signals are $S_A = S_B = [0,1]$. If the state is $\omega = 0$, a pair of signals (s_A, s_B) is drawn in the bottom left triangle $\mathrm{conv}\{(0,0),(1,0),(0,1)\}$ uniformly at random. For $\omega = 1$, one uses the top right triangle $\mathrm{conv}\{(1,1),(1,0),(0,1)\}$; see Figure 1. One can easily check that after observing s_A, Alice's posterior coincides with the signal,[2] i.e., $p_A = s_A$; similarly $p_B = s_B$. The resulting distribution of posteriors (p_A, p_B) is uniform on $[0,1]^2$.

This example motivates the question: Which other product distributions that are symmetric with respect to the state and the agents are feasible, and which are not? [Arieli et al. 2021] give a simple characterization in which the uniform distribution plays a special role. We say that a distribution $\nu \in \Delta([0,1])$ is *symmetric with respect to the state* if $\nu([0,a]) = \nu([1-a,1])$ for every $a \in [0,1]$.

Proposition 1. *When $\nu \in \Delta([0,1])$ is symmetric with respect to the state, $P = \nu \times \nu$ is feasible if and only if the uniform distribution on $[0,1]$ is a mean-preserving spread[3] of ν.*

In other words, Proposition 1 states that the uniform distribution is the most informative among symmetric product distributions. One direction of Proposition 1

[2]By a revelation principle, any feasible distribution can be implemented via a direct information structure, where signals coincide with induced posteriors. We rely on such information structures in all the examples.

[3]A distribution ν_1 is a mean-preserving spread of ν_2 if ν_1 can be obtained from ν_2 by redistributing the mass of each point $x \in [0,1]$ in a way that the mean is equal to x. Formally, there exists a pair of random variables (X_1, X_2) with distributions ν_1 and ν_2, respectively, such that $\mathbb{E}[X_1 \mid X_2] = X_2$.

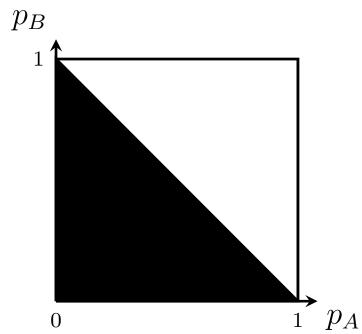

Fig. 1. An implementation of the uniform distribution. The distribution of signals is uniform over the black area for $\omega = 0$ and over the white area for $\omega = 1$.

is simple. [Blackwell 1951] showed that, given an information structure implementing some mean-preserving spread of ν, an information structure implementing ν can be obtained via *garbling*: a post-processing of the signal. Hence, if the uniform distribution is a mean-preserving spread of ν, one can garble the signals in the information structure from Example 3 for each agent independently and induce the distribution $\nu \times \nu$. The opposite direction is more involved.

Example 3 and Proposition 1 indicate that the posteriors of Alice and Bob can be uncorrelated. Can the posteriors be negatively correlated? The following example gives a positive answer and the proposition after it provides a limitation.

Example 4. Consider the distribution $P = \frac{1}{8}\delta_{(\frac{1}{4},1)} + \frac{3}{8}\delta_{(\frac{1}{4},\frac{1}{2})} + \frac{3}{8}\delta_{(\frac{3}{4},\frac{1}{2})} + \frac{1}{8}\delta_{(\frac{3}{4},0)}$ depicted in Figure 2. This example features very counter-intuitive beliefs: Whenever Alice's posterior is $\frac{1}{4}$ (low), her belief about Bob's posterior is $\frac{3}{4}\delta_{\frac{1}{2}} + \frac{1}{4}\delta_1$ (high). Whenever Alice's posterior is $\frac{3}{4}$ (high), her belief about Bob's posterior is $\frac{1}{4}\delta_0 + \frac{3}{4}\delta_{\frac{1}{2}}$ (low). Nevertheless this distribution is feasible. To see this, consider the information structure where at state $\omega = 0$, the pair of signals (s_A, s_B) is either $(\frac{1}{4}, \frac{1}{2})$ or $(\frac{3}{4}, 0)$ with corresponding probabilities $\frac{3}{4}$ and $\frac{1}{4}$. Similarly, for $\omega = 1$, the signals (s_A, s_B) are either $(\frac{3}{4}, \frac{1}{2})$ or $(\frac{1}{4}, 1)$ with probabilities $\frac{3}{4}$ and $\frac{1}{4}$. One can verify that Alice's posterior $p_A = s_A$ and Bob's $p_B = s_B$. Therefore, this information structure implements the distribution P.

The covariance of posteriors for the distribution P in Example 4 is $\mathrm{Cov}(P) = \mathbb{E}_P\left[(p_A - \frac{1}{2})(p_B - \frac{1}{2})\right] = -\frac{1}{32}$. As we show in [Arieli et al. 2021], this covariance is the minimal possible for any joint posterior belief distribution with prior $p = \frac{1}{2}$.

Proposition 2. $\mathrm{Cov}(P) \geq -\frac{1}{32}$ *for any feasible distribution* P.

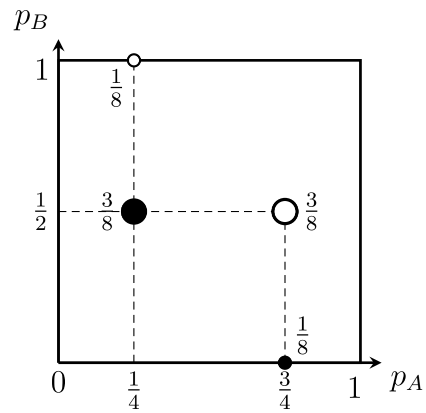

Fig. 2. The distribution P from Example 4 and its implementation. Black points correspond to posteriors in the state $\omega = 0$ and white points to those in $\omega = 1$.

Extreme points of the set of feasible distributions

The set of feasible distributions is a compact convex set,[4] and as such can be naturally represented via its extreme points. When we optimize a linear (or more generally, convex) objective over a compact convex set, this representation becomes especially useful because, by Bauer's principle, the optimum is always attained at an extreme point. In particular, extreme points play an important role in Bayesian persuasion since the sender's objective is linear in the distribution P of beliefs, provided that there are no strategic externalities among the receivers.

In the single-agent case, the extreme points of the set of feasible posterior belief distributions are the binary-support distributions with mean p. As a consequence, two signals are enough for optimal persuasion of one receiver. For more than one agent, no simple characterization of the extreme points is known. The following example from [Arieli et al. 2021] shows the existence of extreme points with countably-infinite support.

Example 5. Figure 3 demonstrates an extreme feasible distribution P together with its implementation. The intuition for extremality of P comes from having no flexibility in its implementation: The mass of $\frac{1}{3}$ at the point $(0, \frac{3}{7})$ must originate from state $\omega = 0$ (because Alice has posterior 0). To have a posterior of $\frac{3}{7}$ for Bob at this point we must assign the entire mass of $\frac{1}{4}$ at the point $(\frac{9}{13}, \frac{3}{7})$ to state $\omega = 1$. We proceed inductively and see that the conditional distributions of posteriors given ω are pinned down uniquely. In [Arieli et al. 2021], we demonstrate that if P was possible to represent as a convex combination of distinct feasible distributions, the implementation would not be unique. Since the implementation is unique, P is extreme.

[4]If P and Q are feasible, then for every $\alpha \in [0, 1]$ the information structure that with probability α reveals information according to P and with probability $(1 - \alpha)$, according to Q, implements $\alpha P + (1 - \alpha)Q$. Compactness requires a more elaborate argument.

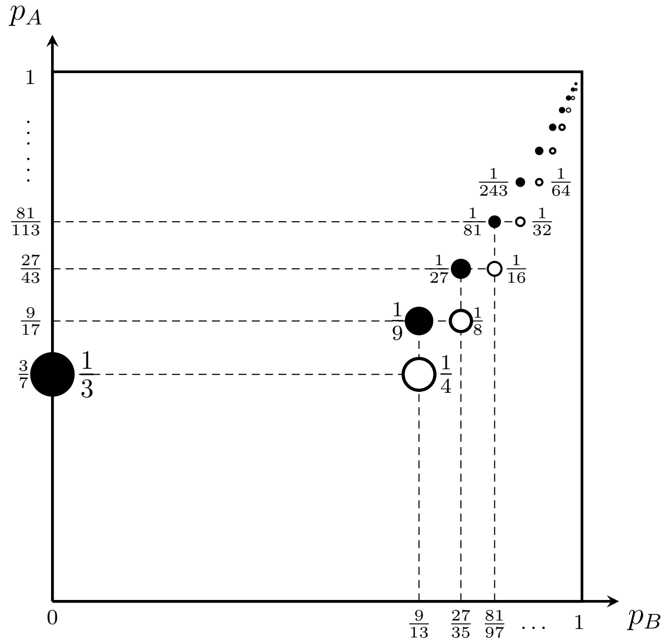

Fig. 3. An extreme feasible distribution P of posteriors with countably-infinite support. The numbers near the points indicate their probabilities. Conditional on $\omega = 0$, the pair of posteriors belongs to the set of black points and, conditional on $\omega = 1$, to the set of white points.

One may ask whether the uniform distribution (Example 3) is an extreme point. The following general result implies a negative answer.

Proposition 3. *Every extreme point of the set of feasible distributions of posteriors is supported on a set having zero Lebesgue measure.*

Consequently, optimal persuasion can always be achieved using only posteriors from a Lebesgue-negligible set.

Feasibility, trade, and agreeing to disagree

So far, we saw several examples of feasible distributions, for which feasibility was shown directly by constructing an information structure implementing the distribution. But how can one prove that a distribution is infeasible? In the case of a finite-support distribution, one can write down a linear program for a direct information structure and check its feasibility. But this program can quickly become very complicated. And what if the distribution has infinite support?

Example 6. Let P be the uniform distribution over $[0,1]^2 \setminus \left([0,0.1]^2 \cup [0.9,1]^2\right)$; see Figure 4. To show infeasibility of P we introduce the notion of a trade.

Assume, by way of contradiction, that this distribution is feasible. Consider a hypothetical scenario where $\omega \in \{0,1\}$ is the value of a good. Alice and Bob each have one copy of this good to trade and their beliefs about the value are distributed according to P. In addition, there is another party, the "mediator", who also has

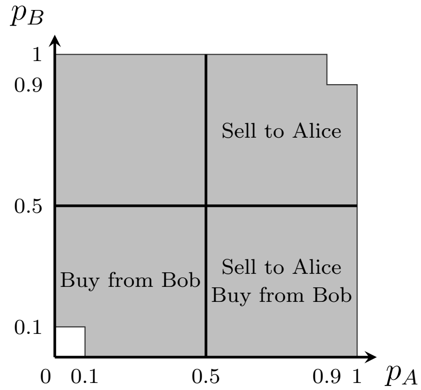

Fig. 4. The distribution P is uniform over the gray area. The outcome of the mediator's trading strategy is indicated in each area.

one copy and observes the private information of Alice and Bob. Consider the following trading strategy for the mediator:

—Whenever Bob's posterior p_B is less than a half, the mediator buys a good from Bob at a price of p_B.

—Whenever Alice's posterior p_A is more than a half, the mediator sells a good to Alice at a price of p_A.

Both Alice and Bob expected gain is zero, since the price coincides with the expected value of the good given the agent's information and the offer itself carries no additional information to what the agent already knows.

Let us evaluate the mediator's expected revenue from this trading strategy.

—The expected payment from the mediator to Bob is $\frac{9}{98} \cdot 0.05 + \frac{40}{98} \cdot 0.3 = \frac{12.45}{98}$.

—The expected payment that the mediator gets from Alice is $\frac{9}{98} \cdot 0.95 + \frac{40}{98} \cdot 0.7 = \frac{36.55}{98}$.

—Whenever the realization of the posteriors is in $[0.5, 1]^2 \setminus [0.9, 1]^2$, the mediator has to supply a good to Alice but he does not get one from Bob. Since the good is worth at most 1, we can bound the mediator's expected losses in this case by its probability, which is $\frac{24}{98}$.

Summing up, we get that the total revenue of the mediator is at least $\frac{36.55-12.45-24}{98} > 0$. However, the mediator's expected gain must equal the agents' total expected loss, which we know is zero. This contradiction implies that P is infeasible.

Consider another example where infeasibility can be proved via a no-trade argument.

Example 7. Let $P = \frac{9}{20}\delta_{(\frac{9}{40},\frac{1}{2})} + \frac{1}{2}\delta_{(\frac{3}{4},\frac{1}{2})} + \frac{1}{40}\delta_{(\frac{19}{40},0)} + \frac{1}{40}\delta_{(\frac{19}{40},1)}$; see Figure 5. The trading strategy of the mediator is as follows:

—Whenever Bob's posterior p_B equals $1/2$ the mediator buys a good from Bob at a price of $1/2$.

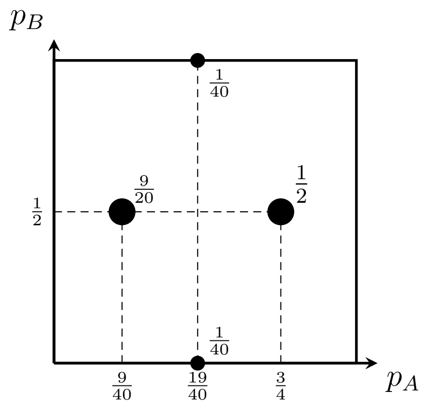

Fig. 5. The distribution P from Example 7.

—Whenever Alice's posterior p_A is in $\{\frac{9}{40}, \frac{3}{4}\}$, the mediator sells a good to Alice at a price of p_A.

A calculation similar to the one above shows that the mediator's expected gain is positive, which shows that P is not a feasible distribution.

In Example 7, the mediator's trading strategy is somewhat sophisticated: the mediator sells the good to Alice when her posterior is $\frac{9}{40}$, but does not do so when the posterior is $\frac{19}{40}$, i.e., higher. In particular, the mediator's strategy does not have a threshold structure — buying whenever the posterior is below some threshold and selling whenever the posterior is above some threshold. One can verify that in this example no threshold trading strategy yields a positive revenue for the mediator.

A main result of [Arieli et al. 2021] is that the above technique can be applied to demonstrate the infeasibility of any infeasible distribution P. The mediator's trading strategy depends on two measurable sets $A, B \subset [0,1]$:

—The mediator buys a good from Bob whenever Bob's posterior is in B, for a price p_B.

—The mediator sells a good to Alice whenever Alice's posterior is in A, for a price p_A.

Denote the marginals of P by $P_1, P_2 \in \Delta([0,1])$. The mediator's expected revenue is bounded from below by

$$\int_A x \,\mathrm{d}P_1(x) - \int_B x \,\mathrm{d}P_2(x) - P(A \times \overline{B}), \tag{1}$$

where the first term is the expected payment of Alice to the mediator, the second term is the mediator's payment to Bob, and the third term bounds the mediator's loss when he supplies his own good to Alice. If the expression in (1) is strictly positive for some A and B, we deduce that P is infeasible, by the same logic used in Examples 6 and 7. We can also swap the roles of Alice and Bob and get a symmetric expression

$$\int_B x \,\mathrm{d}P_2(x) - \int_A x \,\mathrm{d}P_1(x) - P(\overline{A} \times B). \tag{2}$$

Again, the positivity of this expression implies that P is infeasible.

Surprisingly, the requirement that the two expressions in (1) and (2) are non-positive is not only necessary but also sufficient for feasibility of P. Namely, a distribution is feasible if and only if the mediator cannot extract positive expected revenue. This provides a no-trade interpretation for a theorem by Dawid, DeGroot, and Mortera.

Theorem 1 ([Dawid et al. 1995])**.** *A distribution $P \in \Delta([0,1]^2)$ is feasible*[5] *for some p if and only if*

$$P(A \times \overline{B}) \geq \int_A x \,\mathrm{d}P_1(x) - \int_B x \,\mathrm{d}P_2(x) \geq -P(\overline{A} \times B) \tag{3}$$

for all measurable $A, B \subseteq [0,1]$.

In [Arieli et al. 2021], we reinterpret this theorem in the context of Aumann's agreement theorem [Aumann 1976]. Inequality (3) can be seen as a quantitative bound on how much Bayesian agents can agree to disagree. We also show an extension of the no-trade characterization to more than two agents. In the multiple-agent case, more sophisticated trading strategies are required. In particular, to obtain an if-and-only-if characterization, we allow the mediator to trade fractions of the good with the agents. An extension beyond binary states appears in the subsequent paper by [Morris 2020].

REFERENCES

ReferencesRef

Arieli, I., Babichenko, Y., Sandomirskiy, F., and Tamuz, O. 2020. Feasible joint posterior beliefs. In *Proceedings of the 21st ACM Conference on Economics and Computation.* 643–643.

Arieli, I., Babichenko, Y., Sandomirskiy, F., and Tamuz, O. 2021. Feasible joint posterior beliefs. *Journal of Political Economy (to appear).*

Aumann, R. J. 1976. Agreeing to disagree. *The Annals of Statistics*, 1236–1239.

Aumann, R. J. and Maschler, M. 1995. *Repeated games with incomplete information.* MIT press. In collaboration with Richard E. Stearns.

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

Dawid, A., DeGroot, M., and Mortera. 1995. Coherent combination of experts' opinions. *Test 4,* 2, 263–313.

Gutmann, S., Kemperman, J., Reeds, J., and Shepp, L. A. 1991. Existence of probability measures with given marginals. *The Annals of Probability*, 1781–1797.

Kamenica, E. and Gentzkow, M. 2011. Bayesian persuasion. *American Economic Review 101,* 6, 2590–2615.

Morris, S. E. 2020. No trade and feasible joint posterior beliefs. *A working paper*.

[5] [Dawid et al. 1995] use different terminology. Feasible distributions are called there *coherent laws.*